MNRAS **000**, 1–30 (2015) Preprint 22 December 2015 Compiled using MNRAS LATEX style file v3.0

Kinematic Modelling of Disc Galaxies using Graphics Processing Units

G. Bekiaris,[1⋆] K. Glazebrook,[1] C.J. Fluke,[1] R. Abraham[2]
[1] *Centre for Astrophysics and Supercomputing, Swinburne University of Technology, P.O. Box 218, Hawthorn, VIC 3122, Australia*
[2] *Department of Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H8, Canada*

arXiv:1510.01815v2 [astro-ph.IM] 20 Dec 2015

22 December 2015

ABSTRACT

With large-scale Integral Field Spectroscopy (IFS) surveys of thousands of galaxies currently under-way or planned, the astronomical community is in need of methods, techniques and tools that will allow the analysis of huge amounts of data. We focus on the kinematic modelling of disc galaxies and investigate the potential use of massively parallel architectures, such as the Graphics Processing Unit (GPU), as an accelerator for the computationally expensive model-fitting procedure. We review the algorithms involved in model-fitting and evaluate their suitability for GPU implementation. We employ different optimization techniques, including the Levenberg–Marquardt and Nested Sampling algorithms, but also a naive brute-force approach based on Nested Grids. We find that the GPU can accelerate the model-fitting procedure up to a factor of ~ 100 when compared to a single-threaded CPU, and up to a factor of ~ 10 when compared to a multi-threaded dual CPU configuration. Our method's accuracy, precision and robustness are assessed by successfully recovering the kinematic properties of simulated data, and also by verifying the kinematic modelling results of galaxies from the GHASP and DYNAMO surveys as found in the literature. The resulting GBKFIT code is available for download from: `http://supercomputing.swin.edu.au/gbkfit`.

Key words: galaxies: kinematics and dynamics – galaxies: spiral – methods: data analysis

1 INTRODUCTION

Galaxy kinematics play a crucial role in answering many fundamental questions about the structure, formation and evolution of galaxies. The missing mass problem which was first raised by Zwicky (1933) and later confirmed by the flat rotation curves found in the majority of observed disc galaxies (Rogstad & Shostak 1972; Roberts & Rots 1973; Bosma 1978; Rubin et al. 1978), eventually led to the theory of dark matter (Trimble 1987; Ashman 1992) which revolutionised our understanding of galaxies. Modern theories of galaxy formation are based on the ΛCDM theory (Davis et al. 1985; Efstathiou et al. 1990; Planck Collaboration et al. 2015), where galaxies form and grow inside massive dark matter haloes that dominate the baryons dynamically and whose merging history drives galaxy growth (White & Rees 1978; Cooray & Sheth 2002; Mo et al. 2010). Galaxy kinematics play a vital role in studying this dynamical evolution (Kannappan 2001; de Rijcke et al. 2005; Erb et al. 2006; Weiner et al. 2006).

When studying the kinematics of disc galaxies, one of the most important measurements is the rotation curve (Sofue & Rubin 2001). By studying the shape of the rotation curve we can infer or constrain various dynamical quantities such as the baryonic and dark matter mass distribution (Broeils & Courteau 1997; Spano et al. 2008). Rotation curves allow us to derive the maximum circular velocity and thereby construct the Tully-Fisher relation (Tully & Fisher 1977) which is a correlation between the intrinsic luminosity of a galaxy and its rate of rotation. The Tully-Fisher relation can be used to measure distances and peculiar velocities (Mathewson et al. 1992), constrain the structural properties of dark matter haloes (Gnedin et al. 2007) or indicate merger activity (Kannappan et al. 2002).

Integral Field Spectroscopy (IFS) allows us to obtain spatially-resolved spectral information across the sky, enabling us to investigate galaxy kinematics in greater detail. IFS removes many of the problems and limitations of long-slit spectroscopy (Bacon et al. 1995) and reveals all the non-circular and peculiar velocity structure across the observed galaxy while allowing us to derive orientation parameters, including the major kinematic axis, from the observed data.

Galaxy kinematics based on IFS data can be used to study a wide variety of topics including: galaxy for-

⋆ E-mail: gbekiaris@swin.edu.au

© 2015 The Authors

mation and evolution scenarios (Falcón-Barroso et al. 2004; Yang et al. 2008; Genzel et al. 2008; Cappellari et al. 2011; Wisnioski et al. 2011), the supermassive black holes at the centre of galaxies (Verolme et al. 2002) and galaxies with exotic kinematic signatures (Davies et al. 2001).

To extract the rotation curve from a galaxy we create kinematic models and we fit them to observations. For high resolution observations the tilted-ring modelling approach (Rogstad et al. 1974), which extracts the rotation curve in a piecewise manner, has become a standard. This method uses a series of parametrized rings to describe the circular velocities at different radii of the galactic disc. For low resolution observations, where high-parameter models cannot be used due to low signal-to-noise ratio and beam smearing effects, different model are employed. Usually these models combine a parametric function for the rotation curve model, the inclination and orientation of the galaxy, and a photometric profile.

Parametric rotation curve models can have a theoretical basis or be purely empirical. Models such as the exponential disc (Freeman 1970) and the pseudo-isothermal sphere (Begeman 1987) are derived from the stellar and dark matter mass distribution respectively. The linear ramp model (Wright et al. 2007), although it has no theoretical basis, seems to describe high-redshift data very well (Epinat et al. 2010). The more elaborate empirical Multi-Parameter Function of Courteau (1997) can be used to describe rotation curves with an increasing or decreasing velocity plateau.

In order to compare models to observations, one has to take the spatial resolution into account, by convolving the model with the Point Spread Function (PSF) of the observation. While this is not very important for high resolution local galaxy data, it is essential for high redshift galaxies and low resolution observations where the size of the PSF is comparable to the angular size of the galaxy (Förster Schreiber et al. 2006; Gnerucci et al. 2011; Epinat et al. 2012).

Several kinematic model-fitting codes are publicly available. ROTCUR (Begeman 1987) derives the kinematic parameters from the observed velocity field by fitting a tilted-ring model to it. DISKFIT (Spekkens & Sellwood 2007) fits a physically motivated model that handles non-axis-symmetric rotation to a velocity map. KINEMETRY (Krajnović et al. 2006) extracts the kinematic properties of a galaxy by using a generalisation of photometry to higher-order moments of the line-of-sight velocity distribution.

Although most of the methods discussed in the literature (including the above codes) fit their models to 2D projections, such as the velocity map, there are cases where the fit is done directly to the observed intensity datacube. By doing so, one can take into account the effects of beam smearing, including the asymmetries on the spectral lines and the increase of the velocity dispersion towards the centre of the galaxy (Davies et al. 2011). TIRIFIC (Józsa et al. 2007) and ^{3D}BAROLO (Di Teodoro & Fraternali 2015) use a tilted-ring approach to generate a parametrized model datacube and fit it directly to the observation. GALPAK^{3D} (Bouché et al. 2015) extracts the morphological and kinematic parameters of a galaxy by fitting a ten-parameter disc model to three-dimensional data. These, and other model-fitting approaches, employ an optimization algorithm to fit the kinematic model to the observed data.

While the Levenberg–Marquardt algorithm (Levenberg 1944; Marquardt 1963) has been the most popular choice for optimization of parameters in kinematic modelling (e.g., Barbosa et al. 2006; Coccato et al. 2008; Green et al. 2014), other optimization methods have also been used successfully. Some of the most notable include the Nelder-Mead minimizer (e.g., Garrido et al. 2004), Markov Chain Monte Carlo samplers (e.g., Wisnioski et al. 2011; Bouché et al. 2015) and genetic algorithms (e.g., Cresci et al. 2009; Wiegert & English 2014).

The fitting procedure can be a time-demanding process, especially when it involves complex many-parameter models, high resolution data, and optimization algorithms that require a substantial number of iterations before converging to a solution. Unfortunately most existing kinematic fitting software does not fully utilise all the available computer hardware efficiently in order to decrease the computational time requirements.

The lack of high-performance software for galaxy kinematic modelling has not been a major problem so far because of the limited amount of 3D spectroscopic data available for analysis. Some of the largest spatially resolved spectroscopic surveys completed to date include: GHASP (Epinat et al. 2008a), SINS (Förster Schreiber et al. 2009), DiskMass (Bershady et al. 2010), and ATLAS^{3D} (Cappellari et al. 2011). These surveys consist of up to a few hundreds of galaxies, with the last being the largest one (260 galaxies). However, the number of observed galaxies will increase substantially in the near future because of the large-scale optical and radio surveys such as CALIFA (Sánchez et al. 2012), KMOS^{3D}(Wisnioski et al. 2015), SAMI (Bryant et al. 2015), MaNGA (Bundy et al. 2015), HECTOR (Bland-Hawthorn 2015) and WALLABY (Koribalski 2012), which are currently underway or planned. Furthermore, new- and next-generation instruments like MUSE (Bacon et al. 2015) will be able to capture massive amounts of optical spectra in a single exposure, making the analysis of each individual galaxy computationally challenging.

The vast amounts of data that will be produced by the above surveys will require a substantial amount of computational power for their analysis. Therefore, new high-performance tools and software methods that will accelerate the kinematic modelling process are needed. In addition, high-performance software can also benefit the kinematic analysis of numerical simulations where thousands of line-of-sights and time steps can be extracted (e.g., Perret et al. 2014).

Although in some of the above-mentioned codes there has been some effort to accelerate the fitting procedure by utilising all the available cores of the Central Processing Unit (CPU), none is currently able to take advantage of other many-core architectures such as the Graphics Processing Unit (GPU).

Over the past decade the GPU has been transformed from a specialised graphics rendering hardware into a massively parallel architecture with high computational power. Its enormous memory bandwidth and arithmetic capabilities make it highly suitable for scientific computing. GPUs have already been used for different problems in astronomy, including radio astronomy cross-correlation (Wayth et al. 2009), gravitational microlensing (Thompson et al. 2010),

N-Body simulations (Belleman et al. 2008), real-time control of adaptive optics (Gratadour et al. 2014), and astronomical data visualization (Hassan et al. 2013). In most of the above cases the GPUs provided a substantial performance increase compared to the CPU.

In this work we investigate how the graphics processing unit can benefit galaxy kinematic modelling by accelerating the model fitting procedure. In Section 2 we describe the modelling method including the choice of kinematic models, the different optimization techniques used and the GPU implementation of the fitting algorithm. In Section 3 we introduce and provide details about the datasets we have selected for this work. In Section 4 we evaluate the consistency and robustness of the presented method by fitting models to the selected data and comparing our results with the literature. In section 5 we evaluate the performance of the proposed approach and we compare it to CPU-based solutions. In section 6 we discuss and interpret our findings, but also suggest future research and developments that can follow this work. Finally in Section 7 we list a summary of our conclusions.

2 THE METHOD

Fitting models to observations of galaxies in order to extract their kinematic and structural parameters is a very common problem in astronomy. The fitting procedure employs a model with a number of free parameters, the observed data and an optimization technique which chooses a 'better-fitting' solution based on a previous, less-optimal fit. Although the model and optimization technique used for the fitting procedure may vary, in most cases the following pattern is more or less respected:

(1) Evaluate the model on a cube using the current set of parameters.
(2) Convolve the model cube with the line and point spread functions of the observation.
(3) Compare the convolved model cube or its moments with the observed data and evaluate the goodness of fit.
(4) Check termination criteria and, if met, return the optimized model parameters. Otherwise, adjust model parameters and repeat from Step 1.

The above steps describe a typical iteration of the fitting procedure. Steps 1 to 3 usually remain the same regardless of the chosen optimization technique, while step 4 can differ substantially. The number of iterations required to converge to an optimal solution is usually heavily dependent on the chosen optimization technique, the number of free parameters in the selected model, and the choice of initial model parameter values.

The method we present in this paper follows a similar approach to the one described above but employs two different modelling strategies. The first one uses a velocity field model instead of an intensity cube, and resembles the method used in Epinat et al. (2008b). It does not take the line and point spread functions into account and thus it is more appropriate when modelling high-resolution observations. The second one uses an intensity cube model and resembles the method used in Green et al. (2014). It deals with beam smearing effects by taking into account the line and point spread functions, hence it is a much better option when working with low-resolution data.

While our method is very similar to previous work, its implementation differs substantially from the majority of the available kinematic model-fitting codes since it takes advantage of the massively parallel architecture of the GPU. More detailed information on the fitting procedure and its implementation on the GPU can be found in Section 2.2.

2.1 The models

Every galaxy model in this work is a combination of a disc and a rotation curve. In addition, when an intensity cube model is used, a photometric profile and an intrinsic velocity dispersion are employed as well. We now describe these model components in detail.

2.1.1 The disc model

To model the geometry of the galaxy we use a thin and flat disc. To align any given point (x, y) of the observation with the coordinate system defined by the projected major and minor axes of the disc, we use the following equations:

$$x_e = -(x - x_o)\sin PA + (y - y_o)\cos PA, \tag{1}$$

$$y_e = -(x - x_o)\cos PA - (y - y_o)\sin PA, \tag{2}$$

where (x_o, y_o) is the centre and PA the position angle of the projected disc. The radius r of the circle on the disc plane that passes through the projected position (x, y) is calculated as follows:

$$r = \sqrt{x_e^2 + \left(\frac{y_e}{\cos i}\right)^2}, \tag{3}$$

where i is the inclination of the disc. Finally, we calculate the position angle, θ, of the projected point (x, y) relative to the disc major axis using:

$$\cos\theta = \frac{x_e}{r}. \tag{4}$$

2.1.2 The surface brightness model

The surface brightness of a disc galaxy can be modelled using two components: an inner spheroid component for the bulge and an outer exponential component for the disc (de Vaucouleurs 1959).

In this work we only used an exponential disc profile. While this distribution describes correctly the stellar component, it is less appropriate for describing the ionised gas lines whose emission is dominated by HII regions. However, when dealing with low-resolution data, due to the effects of beam smearing, an exponential disc profile can be a reasonable choice (Cresci et al. 2009; Wright et al. 2009; Green et al. 2014). On the contrary, when working with high-resolution data where the HII regions are clearly resolved, more complex models may be used. In many cases, HII regions are organised into ring- or spiral-like structures but they can also be randomly distributed across the galactic disc (Tsvetanov & Petrosian 1995; Evans et al. 1996). This makes their modelling a non-trivial task. Since the focus of

this work is not to investigate modelling in detail, the decision to use only an exponential disc model can be considered acceptable. Furthermore, by using only an exponential disc profile we were able to compare our modelling results directly to the results found in the literature (see Section 4.3) However, our code is flexible enough to cope with a wide range of user-specified intensity models.

The exponential disc profile is described by the following equation:

$$I(r) = I_0 \exp\left(-\frac{r}{r_0}\right), \tag{5}$$

where $I(r)$ is the surface brightness at radius r, I_0 is the surface brightness at the centre of the galaxy, and r_0 is the scale length.

2.1.3 The rotation curve models

In this work we have used two different rotation curve models: an *arctan* profile (Courteau 1997) and the analytical function of Epinat et al. (2008a) which from now on we refer to as the *Epinat function*.

The two-parameter arctan function is given by:

$$V_{\rm rot}(r) = \frac{2}{\pi} V_{\rm t} \arctan\frac{r}{r_{\rm t}}, \tag{6}$$

where $V_{\rm rot}(r)$ is the circular velocity at radius r, $V_{\rm t}$ is the asymptotic circular velocity and $r_{\rm t}$ the turn-over radius. Since $V_{\rm t}$ is never reached in practice, this model describes rotation curves with an increasing velocity plateau.

The four-parameter Epinat function is given by:

$$V_{\rm rot}(r) = V_{\rm t}\frac{(r/r_{\rm t})^g}{1+(r/r_{\rm t})^a}, \tag{7}$$

where $V_{\rm t}$ and $r_{\rm t}$ control the turn-over velocity and radius while g and a control the inner and outer slope of the rotation curve. As a result, this function allows us to describe rotation curves with flat ($a = g$), rising ($a < g$), and declining ($a > g$) velocity plateaus.

2.1.4 The velocity dispersion model

Our galaxy model includes an intrinsic velocity dispersion component, $\sigma_{\rm model}$, which is assumed to be constant across the galactic disc. Fitting a disc model which includes the velocity dispersion allows us to correctly account for beam smearing (Davies et al. 2011), but it can be problematic when the velocity field is not representative of a disc or if the intrinsic velocity dispersion is not constant (Green et al. 2014).

An alternative method for extracting the intrinsic velocity dispersion of a galaxy is to use the $\sigma_{\rm m}$ measure of Law et al. (2009), which is the flux-weighted mean of the velocity dispersions measured in each individual pixel. This quantity provides a more robust way to measure the velocity dispersion when working with data of low signal-to-noise ratio and it does not require assumptions about an underlying kinematic model. However, it can be biased due to beam smearing effects. To correct for this, Green et al. (2014) subtracted a map of unresolved velocity shear present in the best-fitting disc model from the observed velocity dispersion map before calculating $\sigma_{\rm m}$.

Since the purpose of this work is to evaluate the feasibility of the GPU for the kinematic model-fitting procedure, and not to study the velocity dispersion of galaxies in detail, the assumption of a constant intrinsic velocity dispersion is reasonable. Therefore, all the velocity dispersion results derived by our method in this paper refer to the fitted velocity dispersion $\sigma_{\rm model}$.

2.1.5 The complete kinematic model

The observed radial velocity in the sky plane at (x, y) of a purely rotating thin disc with systemic velocity $V_{\rm sys}$ is described by:

$$v_{\rm model}(x, y) = V_{\rm sys} + V_{\rm rot}(r)\sin i\cos\theta. \tag{8}$$

The above radial velocity model has seven or nine free parameters depending on the selected rotation curve profile.

Our final model can be either a velocity field or an intensity cube. In the case of a velocity field, equation 8 is evaluated for each pixel directly. In the case of an intensity cube, equation 8 along with equation 5 and $\sigma_{\rm model}$ are used to infer the shape of the spectral lines prior to the convolution step. However, as a last step, the velocity and velocity dispersion maps are extracted from the intensity cube model. This is done because our fitting method always expects observations of velocity and velocity fields as input. For more details on the evaluation of the velocity field and intensity cube models see Sections 2.2.1 to 2.2.5.

2.1.6 The initial model parameter values

In order to start the fitting procedure, an initial set of model parameter values has to be given to the optimization algorithm. Ideally we want the initial parameter values to be as close to the optimal solution as possible in order to decrease our chances of getting trapped into a local optimum. These values are calculated by analysing the continuum, velocity and velocity dispersion maps of the observation.

The coordinates of the kinematic centre ($x_{\rm o}$,$y_{\rm o}$) are estimated by calculating the central moments of the continuum map. $V_{\rm sys}$ is calculated by extracting the velocity mean around the previously estimated kinematic centre from the velocity map. To estimate PA, we draw a line that passes through the estimated kinematic centre and splits the velocity field into two parts. The orientation of the line is selected in a way that the difference of the means of the two parts is maximized. The initial values of i and $r_{\rm t}$ are set to the photometric inclination and scale radius found in online databases such as HyperLeda (Paturel et al. 2003, for the GHASP survey) and SDSS (Adelman-McCarthy et al. 2006, for the DYNAMO survey). For the $V_{\rm t}$ and $\sigma_{\rm model}$ parameters we use the values of 150 km s^{-1} and 30 km s^{-1} respectively. The a and g parameters of the Epinat rotation curve function are both set to 1.

The parameter values of the surface brightness profile are assumed to be known and they remain fixed during the fitting procedure. The value of I_0 is technically a scaling factor, and since our fit is done to the velocity and velocity dispersion maps, its value is not expected to affect the fitting results. The scale length (r_0) is fixed to its actual value (for the mock data) or to the value found in the SDSS

database (for the DYNAMO survey). This is a reasonable decision since the scale length can be determined from the photometry.

More information about the simulated data and the two surveys studied in this work (i.e., GHASP and DYNAMO) can be found in Section 3.

2.1.7 The point and line spread functions

Finally, we need to know how the imaging system responds to a point source. This is provided by the point and line spread functions. The point spread function (PSF) is modelled using a two-dimensional Gaussian function and its standard deviation, $\sigma_{\rm PSF}$, depends primarily on the atmospheric seeing. The line spread function (LSF) is modelled using a one-dimensional Gaussian function and its standard deviation, $\sigma_{\rm LSF}$, depends on the instrument's spectral resolution. These two functions can be combined to form a three-dimensional Gaussian function:

$$f(x,y,z) = \frac{1}{\sigma_{\rm PSF}^2 \sigma_{\rm LSF} \sqrt{(2\pi)^3}} \exp\left(-\frac{x^2+y^2}{2\sigma_{\rm PSF}^2} - \frac{z^2}{2\sigma_{\rm LSF}^2}\right). \quad (9)$$

2.2 The fitting procedure and implementation

In this section we provide details about our fitting method and how it is implemented on the graphics processing unit.

GPU-programming requires a very different mental model from that required for programming general purpose processors because of their substantially different architectures. CPUs feature a small number of processing cores, large caches, and units for managing and scheduling tasks such as branch prediction, instruction ordering and speculative execution. On the other hand, GPUs sacrifice the large caches and the advanced management and scheduling units in order to maximize the number of processing cores on the chip (Thomas et al. 2009).

Previous research on achieving high performance on a GPU has already been done (Ryoo et al. 2008; Barsdell et al. 2010; Brodtkorb et al. 2013), and although each application might require manual tuning, the following are considered rules of thumb:

- **Maximize parallelism:** In order to fully utilise the GPU one has to make sure that all its processing cores are occupied at any given time. One way to achieve this is by using data-parallel algorithms (Hillis & Steele Jr 1986) which perform simultaneous operations across large sets of data using the parallel processors of the underlying platform. However, the performance of such algorithms may suffer if the size of the data is small relative to the number of processing cores available on the GPU.
- **Maximize arithmetic intensity:** Arithmetic intensity is defined as the number of mathematical operations performed per memory access. Although GPUs are coupled with high-bandwidth main memory (also known as *global memory*) and a small but very fast on-chip memory, their strongest point is to execute arithmetic instructions. As a result, if a GPU program performs only a few arithmetic operations per memory access it becomes *memory-bound* and its performance is limited by the rate its memory can be accessed.
- **Minimize time spent on memory accesses:** Instead of (or in addition to) increasing the arithmetic intensity of the program by increasing the number of instructions per memory access, one can minimize the time spent on global memory transactions. One way to achieve this is to move frequently-accessed data and data that remain constant during the code execution to the on-chip shared memory and the constant memory respectively. The advantage of shared and constant memory over global memory is the superior bandwidth and the much lower latency. Another strategy to reduce the time spent on memory transactions is to use coalesced global memory access by making sure that successive GPU threads access adjacent locations in a contiguous stretch of memory. This strategy takes advantage of the fact that whenever a GPU thread reads or writes the global memory it always accesses a large chunk of memory at once. If other threads make similar accesses at the same time then the GPU can reuse the same chunk for those threads without attempting a new memory transaction.

Due to the image nature of our data we expect our method to benefit from GPU parallelization. 3D spectroscopy data involves 2D and 3D images of thousands or millions of pixels that can be processed independently and thus we expect a high GPU utilisation. Furthermore, due to the well behaved and regular structure of image data it is relatively easy to write algorithms that follow particular memory access patterns and achieve coalesced global memory access.

To investigate whether our method is suitable for GPU parallelization we follow a similar approach to the one discussed in Barsdell et al. (2010). We outline each step involved in the fitting iteration and then identify which ones resemble known algorithms. For each identified algorithm we refer to its pre-existing analysis, while for the steps that do not appear to resemble any known algorithms we perform the analysis ourselves.

Our method fits a model to the velocity and velocity dispersion maps of the observation. Because of this, when we are using an intensity cube model the method extracts the velocity and velocity dispersion maps from the convolved datacube in order to compare them to the input data and calculate the goodness-of-fit. In addition, one can also supply measurement error and mask maps if they are available. In Sections 2.2.1 to 2.2.7 we describe all the steps involved in the two modelling strategies in detail.

2.2.1 Velocity field evaluation

For the velocity field modelling, the first step of the fitting procedure evaluates the radial velocity model (equation 8) on a two-dimensional grid using the current set of model parameters. The evaluation of the velocity field is a highly parallel transform operation because each pixel evaluation is completely independent of the others. To implement the velocity field evaluation on the GPU we assign each thread to evaluate a different pixel. When working with low-resolution data the number of pixels in our velocity field might be less than the number of available GPU cores and thus the GPU may remain underutilised. Therefore the velocity field modelling strategy will perform best with high-resolution data. In order to achieve high performance we also make sure that

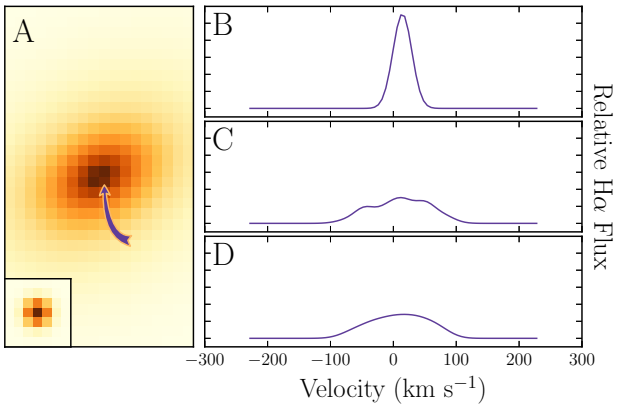

Figure 1. Artefacts may appear on the spectral lines of low resolution model datacubes after the PSF convolution. (A) Integrated flux of a low resolution model datacube before the convolution with the PSF. (B,C,D) All three plots refer to the same spectral line on datacube A. (B) The spectral line before the PSF convolution. (C) The spectral line after the PSF convolution. (D) The spectral line after the PSF convolution with a high resolution version of datacube A. The spatial and spectral resolutions of the cubes are: 0.7 arcsec pixel^{-1} and 30 km s^{-1} for the original cube, and 0.35 arcsec pixel^{-1} and 15 km s^{-1} for the high-resolution cube.

memory coalescing is achieved by assigning consecutive pixels to consecutive GPU threads. When using the velocity field model strategy the Steps 2.2.2 to 2.2.5 are skipped and the iteration continues at Step 2.2.6.

2.2.2 *Intensity cube evaluation*

For the intensity cube modelling strategy, the first step of the fitting procedure generates an intensity cube defined by two spatial and one velocity coordinate. Using the current set of model parameter values, for each spatial coordinate of the cube we calculate a line-of-sight velocity and velocity dispersion by evaluating the full kinematic model described by equation 8. The spatial intensities of the cube are calculated with an exponential disc profile (equation 5) for which the scale length is assumed to be known in advance. The intensities along the velocity dimension of the cube are modelled using a Gaussian profile with its mean and standard deviation corresponding to the calculated line-of-sight velocity and velocity dispersion respectively.

We define the number of pixels on the spatial plane of the observed data as $N_{\mathrm{DATA,x}} \times N_{\mathrm{DATA,y}}$. Furthermore, we define the number of pixels required to cover the cube's modelled spectral lines and their velocity shift as $N_{\mathrm{DATA,z}}$. The number of pixels along the spatial and velocity dimensions of the intensity cube must be an integer multiple of the number of pixels in the observed data. From now on we will refer to this number as the *upsampling factor* or u. The higher the upsampling factor, the higher the resolution of the intensity cube. When dealing with low resolution observations, evaluating the intensity cube at a higher resolution allows us to avoid artefacts that might appear in its spectral lines after the PSF convolution (Fig. 1).

Padding must be applied to the model cube prior to PSF convolution. First, the cube is padded by $N_{\mathrm{PSF,n}} - 1$ across all its dimensions, where $N_{\mathrm{PSF,n}}$ is the size of dimension n of the PSF image. This is required in order to perform the convolution in the frequency domain through the Fast Fourier Transform (FFT, Duhamel & Vetterli 1990). Then, each dimension of the cube is padded to the next power of two because the majority of the FFT algorithms perform optimally when working with that size (Brandon et al. 2008). The length of each dimension of the cube is:

$$N_{\mathrm{CUBE,n}} = 2^{\lceil \log_2 [u \times (N_{\mathrm{DATA,n}} + N_{\mathrm{PSF,n}} - 1)] \rceil} \quad (10)$$

The evaluation of the cube is a highly parallel transform operation because each pixel evaluation is completely independent of the others. To implement the cube evaluation on the GPU we assign each thread to evaluate a different pixel. Since the number of pixels in a cube is usually more than the available GPU cores, each GPU thread evaluates more than one pixel. In order to achieve high performance we also make sure that memory coalescing is achieved by assigning consecutive pixels to consecutive GPU threads.

2.2.3 *Point spread function convolution*

The PSF of the observation is a three-dimensional image, which from now on we will refer to as the *PSF kernel*, and it is created by combining the two-dimensional image of the point spread function and the one-dimensional image of the line spread function. Alternatively, since in this work both the point and the line spread functions are modelled using Gaussian profiles, the PSF kernel can be generated by evaluating equation (9) directly on a three-dimensional grid.

In this step of the fitting iteration we convolve the generated intensity cube with the PSF kernel of the observation. The naive 3D convolution algorithm, although simple, has a high computational complexity of $O(N^6)$ and hence it is preferred to perform the convolution in the frequency domain where the computational complexity is reduced to $O(N^3 log_2^3 N)$. According to the circular convolution theorem, in order to perform a convolution via FFT, the two input sequences must have the same size. Hence, we need to expand the PSF kernel to the intensity cube size.

The intensity cube and PSF kernel are transformed to frequency space using the FFT which is a highly parallel divide-and-conquer algorithm and suitable for GPU parallelization. Next, a point-wise complex multiplication between the two transformed images, which is essentially a transform operation, is performed on the GPU. Because the data of both the cube and the PSF kernel are located in contiguous memory locations coalesced access is achieved once again. Finally an inverse FFT is applied to the results of the multiplication in order to get the final result of the convolution back in geometric space.

2.2.4 *Down-sampling to the resolution of the observation*

Before proceeding further we have to down-sample the intensity cube in order to create a new cube that matches the resolution of the observed data. This is achieved by a transform and reduce operation on the GPU: We assign each GPU thread to calculate the value of one or more pixels of the down-sampled cube by averaging the corresponding pixels of the original cube. Here we achieve coalesced access to

the down-sampled cube but not for the original cube. As a result we do not expect a high efficiency from this step.

2.2.5 *Velocity and velocity dispersion map extraction*

In this step we extract the velocity and velocity dispersion maps from the down-sampled cube, in order to allow for comparison with the input data. This can be done by fitting an analytic profile (e.g., Gaussian) to each of the cube's spectral lines. Alternatively, one can calculate the first- and second-order moments of each spectrum, which correspond to the velocity and velocity dispersion respectively.

Stone (1989) and Anthony & Granick (2009) investigated the two methods thoroughly and assessed their accuracy, computational requirements and robustness. The first method can achieve superior accuracy compared to the second but it is computationally more expensive and may occasionally fail to converge, returning infinite or undefined results for some parameters. The moment-based method, although computationally faster, it can be heavily biased if the underlying continuum is asymmetric, contains multiple spectral lines, or the data have a low signal-to-noise ratio.

Both methods have successfully been used for galaxy kinematics in the literature (Förster Schreiber et al. 2006; Epinat et al. 2008a). However, the moment-based approach is usually favoured when computation time is an issue (Green et al. 2014). Furthermore, in cases where a continuum is clearly present in the data, more sophisticated methods and techniques are employed, such as template fitting to remove the continuum and then Gaussian line fitting to extract the emission lines, e.g. using PPXF (Cappellari & Emsellem 2004) and GANDALF (Sarzi et al. 2006).

In this work we followed the moment-based approach which is a transform and reduce operation. Each GPU thread is assigned to calculate the first- and second-order moments of one or more spectra. Because consecutive GPU threads process consecutive spectra and since the cube data have a slice-major alignment, memory coalesced access is achieved once again. For additional information as to why the moment-based approach was preferred over Gaussian fitting, see Section 5.2.

2.2.6 *Residual and goodness-of-fit evaluation*

After obtaining the velocity and velocity dispersion model maps from the previous step we compare them to the input data and we evaluate the goodness-of-fit using a statistical test. The majority of the statistical tests require the residual between the observed and the modelled data, which is the pixel-wise difference between the two datasets weighted by the error associated with each measurement. This is a trivial transform and reduce operation which is performed in parallel on the GPU, producing a residual velocity map and a residual velocity dispersion map. Although coalesced memory access is achieved, this step is expected to be memory bound since it has a very low arithmetic intensity and it accesses the memory four times for every pixel in the output residual maps: three reads from the model, input and error maps, and one write to the residual map.

The goodness-of-fit calculation takes place on the CPU instead of the GPU. This was a software design decision we had to take in order to make the implementation of our method optimizer-agnostic. To calculate the goodness-of-fit we first have to transfer the residual maps from the GPU to the CPU memory. Although the size of the residual maps can be considered small, the transfer is done via the low-bandwidth PCI Express bus (PCIe, Budruk et al. 2004). Therefore, this transfer might take up a big portion of the total iteration time if the fitted model involves a small amount of operations.

To evaluate the goodness of fit we use the following quantity (Green et al. 2014):

$$\chi^2 = \sum_{\mathrm{i}=1}^{\mathrm{N}} \left[\frac{(v_{\mathrm{i,model}} - v_{\mathrm{i,obs}})^2}{E(v_{\mathrm{i,obs}})^2} + W \frac{(\sigma_{\mathrm{i,model}} - \sigma_{\mathrm{i,obs}})^2}{E(\sigma_{\mathrm{i,obs}})^2} \right], \tag{11}$$

where N is the number of pixels in the data, v and σ are the velocity and velocity dispersion, E is the measurement error and W is the importance of the velocity dispersion in the goodness-of-fit. Since W is a constant, we incorporate it in the measurement error of the velocity dispersion prior the fitting procedure. In case a velocity dispersion field is not present in the fitting procedure the second term of equation 11 is ignored.

The value of W is decided based on how well we expect the selected velocity dispersion model to describe the observed data. By using a small value, $W < 1$, we decrease the contribution of the velocity dispersion residual to the χ^2 statistic, and we can still obtain high-quality velocity field fits, even if our velocity dispersion model does not describe the observed data very well. Green et al. (2014) used $W = 0.2$ to ensure a good fit to the velocity fit, while they used the intensity weighted mean of the velocity dispersion (σ_{m}, Law et al. 2009) to estimate the velocity dispersion of the galaxy.

2.2.7 *Optimization step*

In this step the selected optimization method decides if the fitting procedure should continue or terminate. The termination criteria depend on the optimization algorithm, but they are usually related to the goodness of fit. If the termination criteria are not met the optimization technique will adjust the model parameter values and a new fitting iteration will take place.

For our implementation we used a modular software design which enabled us to use third party optimization libraries and switch between them during runtime. This allowed us to test our approach using different optimization methods without having to modify the source code each time. While this feature is not common in existing kinematic modelling software, it can be very useful since it enables us to cross-check our fitting results using the same software.

All the optimization libraries we have chosen for our software are CPU-based, thus their logic does not take advantage of the GPU parallelization. This limitation is not an important issue since the computational requirements of the optimization logic are expected to be negligible compared to the computational requirements of the rest of the steps involved in the fitting procedure. However, the above statement might be false if the fitted model is simple and has very small runtime requirements. For more details on how

the runtime requirements of the optimizer's logic compare to the total runtime of the fitting procedure, see Section 5.

2.3 The optimization algorithms

The optimization technique is one of the most important factors in the fitting procedure. It has the main impact on both the goodness-of-fit selected and the computational runtime. In this work we use three different optimization methods: the Levenberg–Marquardt Algorithm (LMA), Nested Sampling (NS) and brute-force optimization based on Nested Grids (NG). The three optimizers were chosen based on the approach they follow to explore the parameter space. This also translates to the number of model evaluations required before converging to a solution, and subsequently to the overall runtime requirements of the fitting procedure. For a brief discussion as to why Markov Chain Monte Carlo (MCMC) methods were not considered as practical solution in this work, see Appendix A.

2.3.1 The Levenberg–Marquardt Algorithm

The Levenberg–Marquardt Algorithm is one of the most commonly used optimization methods for galaxy kinematic modelling in astronomy (e.g., Barbosa et al. 2006; Coccato et al. 2008; Green et al. 2014). Starting from an initial set of parameter values, LMA uses gradient and curvature information to explore the parameter space. Compared to the other techniques, it requires a small number of model evaluations before converging to a solution. LMA is a local optimizer and will converge to the closest minimum to the initial guess. Thus, in order to find the global minimum and avoid getting trapped into a local one, the starting point of the optimization must be as close to the global minimum as possible.

In this work we have used the C version of the MPFIT[1] library (Markwardt 2009) which implements both the 'vanilla' LMA algorithm and a modified version that supports constrained optimization through box constraints. Although LMA provides a few opportunities for parallelization (Cao et al. 2009), the MPFIT library is not parallelized.

2.3.2 Nested Sampling

The Nested Sampling algorithm is a Monte Carlo technique for efficient Bayesian evidence calculation and parameter estimation (Skilling 2004). The algorithm starts with a population of N *active* points drawn from a prior distribution. Then, the point with the lowest likelihood $\mathcal{L}_i$ is recorded and replaced by a point drawn from the prior subject to the constraint that the new point has a likelihood $\mathcal{L} > \mathcal{L}_i$. This procedure is repeated and the population of active points moves progressively to areas of higher likelihood.

Several extensions to the original algorithm have been introduced to increase its efficiency and robustness. To improve the acceptance rate when drawing samples from the prior with the hard constraint $\mathcal{L} > \mathcal{L}_i$, Mukherjee et al. (2006) introduced Ellipsoidal Nested Sampling. Their algorithm uses an elliptical bound around the current set of live points in order to restrict the region from which new samples are drawn. To improve the sampling efficiency for multimodal posteriors, Shaw et al. (2007) introduced Clustered Nested Sampling which allows the formation of multiple ellipsoids centered on the modes of the posterior, and hence reducing the region of the prior sampled. This algorithm was further improved by Feroz & Hobson (2008) and Feroz et al. (2009) to efficiently sample from posteriors with pronounced degeneracies between parameters. Finally, Cameron & Pettitt (2014) introduced Importance Nested Sampling which uses the discarded draws from the NS process (i.e., the sampled points failing to satisfy the constraint $\mathcal{L} > \mathcal{L}_i$) and calculates the Bayesian evidence at up to an order of magnitude higher accuracy than the 'vanilla' NS.

In this work we used the MULTINEST[2] library which implements the NS algorithm with all the extensions mentioned above. MULTINEST is relatively easy to use since NS does not need an initial set of parameter values but rather the user has to define a prior, which in many problems is just a uniform distribution. In addition, its behaviour is controlled by only three main parameters: the number of active points N, the maximum efficiency e, and the tolerance tol.

2.3.3 The Nested Grids Optimizer

The Nested Grid optimizer follows a brute-force approach in order to find the optimal solution to the problem. Initially a multi-dimensional regular grid representing the parameter space is defined. For each node of the grid a model is generated and a goodness-of-fit statistic is derived from it. Then, the node with the best goodness-of-fit is selected and a new grid is built around that area. The newly created grid has the same number of nodes and dimensions as the previous one, but it covers a smaller volume on the parameter space. This procedure is repeated until the area of the parameter space under the grid is smaller than a user-defined accuracy threshold a. The resolution of the grid has to be high enough to detect the features in the parameter space and not miss the global optimum.

Although the Nested Grid optimizer is a brute-force method and it is expected to be computationally inefficient, previous research (Fluke et al. 2011) has shown that, under certain circumstances, brute-force approaches can be acceptable when taking advantage of the enormous computational power of the GPU. Furthermore they are also conceptually easier to implement on GPUs as a first solution.

3 THE DATA

In order to assess our method we tested our software using mock data. This allowed us to specify the exact configuration of the observations, including the spatial and spectral resolution (i.e., the point and line spread functions), the galaxy's kinematic and morphological properties, and the amount of noise in the data. Since the simulated galaxies are generated with known model parameter values, we can use mock data to evaluate the correctness, accuracy and robustness of our fitting method.

[1] http://purl.com/net/mpfit

[2] http://ccpforge.cse.rl.ac.uk/gf/project/multinest

Furthermore, we also investigated how our method behaves under real science scenarios by using data from the GHASP (Epinat et al. 2008a) and DYNAMO (Green et al. 2014) surveys. The two surveys were chosen based on their spatial and spectral resolution. GHASP was selected as a representative of high-resolution survey, while DYNAMO is a good example of a survey with low spatial resolution where correct PSF treatment is essential. Other examples of low-spatial resolution surveys are SINS (Förster Schreiber et al. 2006) and MASSIV (Contini et al. 2012). Summary information about all the data used in this work can be found in Table 1.

3.1 The mock data

Each mock observation consists of an intensity cube convolved with a 3D Gaussian PSF kernel (equation 9). Galaxies are modelled using the arctan rotation curve profile (equation 6). A velocity map and a velocity dispersion map are then extracted from the convolved intensity cube by calculating the first- and second-order moments of the spectral lines respectively. In addition, an intensity map is extracted by integrating the intensity across each spectral line. Finally, a mask image is defined in order to simulate pixels with low signal-to-noise ratio and instrument faults (e.g., dead detector pixels or IFU fibres).

For each velocity and velocity dispersion map we created four distinct versions, each with a different level of noise. For the velocity map we used Gaussian noise of 0, 1, 5, and 20 km s^{-1}. The noise on the velocity dispersion map, since it is a second-order moment, was chosen to be $\sim \sqrt{2}$ times larger[3], resulting noise levels of 0, 1.4, 7, and 28 km s^{-1}. The zero-noise maps (which we will also refer to as the *clean dataset*) can be used to verify the accuracy of our method, while the last three can be used to evaluate how our method behaves under low- and high-noise level scenarios. Gaussian noise is not always realistic since real data can contain correlated noise between adjacent pixels (e.g., due to non-circular motions or sky line residual artefacts). However, the inclusion of such noise would not affect our timing results significantly as it does not change the conceptual structure of the problem, and hence it was not explored further.

We generated 1000 mock datacubes by randomly sampling the parameter space defined by the resolution of the observation, and the galaxy's morphological and kinematic parameters (Table 2). To reduce the degrees of freedom of this parameter space we used a fixed point spread function size, while the turn-over radius r_t and scale length r_0 were calculated as a function of the point spread function's FWHM. In addition, we assumed that the systemic velocity V_{sys} has already been subtracted from the moment maps and thus it is equal to zero for all galaxies. This is a reasonable decision since the value of the systemic velocity is not expected to affect our fitting results.

[3] This mathematically motivated choice was validated by fitting Gaussian profiles to simulated spectra with different signal-to-noise ratios, showing that it is a good approximation in most circumstances.

3.2 The GHASP data

The Gassendi Hα survey of SPirals (GHASP, Epinat et al. 2008a) consists of 3D Hα data of 203 spiral and irregular local ($z \sim 0$) galaxies covering a wide range of luminosities and morphological types.

The observations were made with a Fabry-Pérot interferometer attached at the Cassegrain focus of the 1.93 m telescope at the Observatoire de Haute-Provence in France. Over the course of the observations two different detectors were used: (1) a Thomson IPCS with 256 × 256 pixels of 0.96 arcsec coverage each resulting in a field of view of 4 × 4 arcmin; (2) a GaAs IPCS with 512 × 512 pixels of 0.68 arcsec coverage each resulting in a field of view of 5.8 × 5.8 arcmin. The spectral resolution of the instrument is 30 km s^{-1} while the spatial resolution is limited by the seeing which is $\sim$ 3 arcsec and corresponds to 0.02 – 1.4 kpc depending on the redshift of the observed object.

Adaptive spatial binning techniques based on 2D Voronoi tessellations were applied to the datacubes in order to optimize the spatial resolution to the signal-to-noise ratio. From the resulting cubes a series of two-dimensional images was extracted including radial velocity maps, velocity dispersion maps and continuum maps.

We have selected a subsample of the GHASP survey consisting of 173 galaxies. The other 30 galaxies were not considered in this work because their kinematic analysis results were not available.

3.3 The DYNAMO data

The DYnamics of Newly Assembled Massive Objects (DYNAMO, Green et al. 2014) survey consists of spatially resolved local ($z \sim 0.1$) galaxies with high star formation rates (SFR), selected primarily by Hα luminosity. The main aim of this survey is to study galaxy evolution through comparison with high-redshift galaxies.

The IFS observations of the DYNAMO survey were obtained using different instruments on different telescopes, with the largest sample (54 galaxies) being the one coming from the SPIRAL Integral Field Unit (Kenworthy et al. 2001) attached on the 3.9 m Anglo-Australian telescope, located at the Siding Spring Observatory in Australia. The SPIRAL IFU is a lenslet array of 32 × 16 pixels with each pixel having 0.7 arcsec coverage resulting in a field of view of 22.4 × 11.2 arcsec. SPIRAL was connected to the AAOmega spectrograph (Sharp et al. 2006) which was configured to provide a coverage of 500Å and a resolving power of R $\sim$ 12000. The spatial resolution of the observations was limited by the seeing with a mean of 1.4 arcsec which corresponds to 1.5 – 3.6 kpc depending on the redshift of the observed object and it is comparable to high-redshift samples observed with adaptive optics. The spectral resolution of the instrument was $\sim$ 30 km s^{-1}.

Available DYNAMO data include velocity, velocity dispersion and mask maps. For the first two quantities, measurement error maps are also provided.

For this work we have selected a subset of 44 galaxies observed with the SPIRAL instrument and studied in Green et al. (2014). We discarded galaxies with complex kinematics (i.e., galaxies whose both velocity and velocity

Table 1. Information about the different datasets described in Section 3. We used four versions of the same mock data, each version has a different level of noise: 0, 1, 5, and 20 km s^{-1} on the velocity maps, and 0, 1.4, 7, and 28 km s^{-1} on the velocity dispersion maps.

Survey name	Number of galaxies	Redshift	Number of pixels	Spatial sampling (arcsec pixel^{-1})	Spectral resolution (km s^{-1})	Seeing (arcsec)
Mock data	1000	–	32^2 – 96^2	1.0	10 – 60	4.0
GHASP	173	~ 0	256^2 or 512^2	0.96 or 0.68	~ 30	3.0 (mean)
DYNAMO	44	~ 0.1	16×32	0.7	~ 30	1.4 (mean)

Table 2. Parameter space used to generate the mock dataset.

Parameter	Units	Range (min)	Range (max)
$N_{\rm data,xy}$	pixels	32	96
$PSF_{\rm FWHM}$	pixels	4.0 (fixed)	4.0 (fixed)
$LSF_{\rm FWHM}$	km s^{-1}	20	60
$V_{\rm sys}$	km s^{-1}	0 (fixed)	0 (fixed)
$x_{\rm o}$	pixels	$0.4 \times N_{\rm data,xy}$	$0.6 \times N_{\rm data,xy}$
$y_{\rm o}$	pixels	$0.4 \times N_{\rm data,xy}$	$0.6 \times N_{\rm data,xy}$
PA	degrees	0	45
i	degrees	5	85
$r_{\rm t}$	pixels	$0.5 \times PSF_{\rm FWHM}$	$2.0 \times PSF_{\rm FWHM}$
r_0	pixels	$1.0 \times PSF_{\rm FWHM}$	$3.0 \times PSF_{\rm FWHM}$
$V_{\rm t}$	km s^{-1}	50	300
$\sigma_{\rm model}$	km s^{-1}	20	100

dispersion vary significantly from regular discs) for which kinematic analysis results were not available.

4 FITTING RESULTS

To examine the consistency and robustness of our method we fit the three different kinds of data which are described in Section 3. When fitting data from the GHASP and DYNAMO surveys our method and metrics are adjusted to match the ones found in the literature (Epinat et al. 2008b and Green et al. 2014 respectively) as closely as possible. However, some details are not available in the literature, such as the initial model parameter values and the configuration of the optimization algorithms, therefore some difference between the results is expected.

To verify the results, for each observation we run the fitting procedure multiple times using different optimization methods. For all three datasets we use similar optimizer configurations. The LMA optimizer is adjusted to use a finite-differences Jacobian with a step size of ~ 0.1 for all model parameters. The NS optimizer is configured for $N = 50$ active points, a maximum efficiency $e = 1.0$ and a tolerance $tol = 0.5$. For the NG optimizer we choose an accuracy threshold $a = 0.01$ for $r_{\rm t}$, and $a = 0.1$ for the rest of the model parameters. The number of iterations needed to fit using the NG optimizer with all the model parameters free can be quite high and result in prohibitive runtime requirements. To avoid this, we fix the parameters $V_{\rm sys}$, $x_{\rm o}$, $y_{\rm o}$ and i to their actual values (for the mock data) or to a previous best fit (derived by using another optimizer). The grid resolution of the NG optimizer is 7, 11, 11, 5, 5, 5 for the parameters PA, $r_{\rm t}$, $V_{\rm t}$, $\sigma_{\rm model}$, a, and g respectively.

4.1 Fitting the mock data

First, the kinematic model is fitted to the mock data with zero noise. This allows us to evaluate the accuracy of our method but also investigate the robustness of the three optimization techniques. Due to the absence of noise in the data, the method for calculating the initial values for the parameters (Section 2.1.6) may find the best fit for $V_{\rm sys}$, $x_{\rm o}$, $y_{\rm o}$, and PA immediately. When fitting real data, this is rarely the case since the observations include noise and non-circular motions. To account for this, after estimating the initial values for the above parameters, we modified them by a random amount. Also, we assume no prior knowledge about the galaxy's inclination. Therefore, the initial guess for the inclination parameter is set to 45°and its prior constraints between 1°and 89°. During the goodness-of-fit calculation step of the fitting procedure (Section 2.2.6), the dispersion map is taken into account using equation 11 with an importance of $W = 1$.

The fitting results for the clean dataset using the LMA optimizer are shown in the first row of Fig. 2A, B, C, and D. Overall, the accuracy of our method with all three optimizers is excellent: The residuals between the actual and fitted parameters for $\gtrsim 98$ per cent of the galaxies are less than a reasonable threshold (Table 3). In addition, the median residual of each parameter is well below that threshold by 1 – 3 orders of magnitude (Table 4). The few inconsistencies seen are mostly for galaxies with high inclinations.

Low-resolution mock observations of nearly edge-on galaxies, created using a thin and flat disc model, can result in a very small number of pixels in the velocity field. This, in combination with the PSF convolution, can introduce new local minima in the χ^2 landscape and eventually lead the LMA to a wrong solution. In any case, when modelling close to edge-on galaxies, the thickness of the galaxy disc plays an important role, therefore a thin disc model is not a good representative of the observation. Furthermore, the NG optimizer can also converge to a wrong solution if the grid resolution is not sufficient to capture the gradient of the posterior modes. However, we found that we can achieve a ~ 100 per cent success rate if we re-run the fitting procedure using different optimizer settings or initial model parameter values/constraints.

To assess how our method behaves with noisy observations our kinematic model is fitted to all three noisy versions of the mock data. This allows us to investigate how different levels of noise can affect the final fitting results and if the kinematic properties of the galaxies are actually recoverable. However, as it has already been mentioned in Section 3.1, the Gaussian noise in the noisy datasets does not give any indication on the level of non-circular motions that may also

prevent the model to describe correctly the rotating component.

The fitting results for the noisy datasets using the LMA optimizer are shown in the second, third, and fourth rows of Fig. 2A, B, C, and D. Estimating the inclination and/or the circular velocity of nearly face-on galaxies can be quite problematic due to the degeneracy between the two (Begeman 1987). Indeed, a scatter is observed for galaxies with $i \lesssim 20^{\circ}$ (Fig. 2C, empty circles), which subsequently causes a scatter in the circular velocity (Fig. 2D). Overall, the parameters of the majority of the galaxies are recovered with a reasonable accuracy: ~ 68 and ~ 95 per cent of the fitted parameters have a standardised residual within 1σ and 2σ confidence intervals respectively (Table 5).

Further discussion on the fitting results of the mock dataset as well as the degeneracy between the galaxy inclination and circular velocity can be found in Section 6.

4.2 Fitting the GHASP data

To fit the GHASP data we employed the velocity field modelling strategy (Section 2.2.1) and the Epinat rotation curve profile (equation 7), resulting in a nine-parameter problem. However, during the fitting procedure we adopted the kinematic centre used in Epinat et al. (2008b), which reduces the number of free parameters in our model to seven. Furthermore, in order to achieve a good fit for some observations with low signal-to-noise ratio or spatial coverage, the position angle and/or the inclination were fixed to their morphological values.

The model parameter errors returned by the optimizer are unrealistically small. To obtain more realistic errors we use a Monte Carlo method which takes into account the residual velocity field resulting from the fitting procedure. The residual velocity field has a non-uniform structure which is mainly caused by non-circular motions (e.g., bars, inflows, outflows, disc warps etc.). We simulate 200 residual velocity fields from the real velocity field by randomly shuffling the positions of its pixels. Then, we generate 200 simulated velocity fields by adding the simulated residuals to the velocity field of the best fit. Finally, we compute the standard deviation of the best-fitting parameters over the 200 simulated velocity fields.

In Fig. 3 we plot the residuals between our best fits and the best fits of Epinat et al. (2008b). For most of the galaxies there is a good agreement between the results. The median of the residuals for the fitted parameters are: ~ 2 km s^{-1} for the systemic velocity, $\sim 1^{\circ}$ for the position angle, $\sim 2^{\circ}$ for the inclination, and ~ 8 km s^{-1} for the maximum circular velocity. Nonetheless, a few inconsistencies between the results are observed. This is not surprising since the optimizer configuration and priors used in this work are most likely different from the ones used in Epinat et al. (2008b). Furthermore, our fitting procedure was completely automated and we did not treat any galaxy differently. For a complete list of the fitting results see Table C1.

4.3 Fitting the DYNAMO data

To fit the DYNAMO data we employed the intensity cube modelling strategy (Section 2.2.2) and the arctan rotation curve profile (equation 6), resulting in an eight-parameter problem. Due to the low spatial resolution of the data and the degeneracy between the circular velocity and disc inclination, the latter is kept fixed to the photometric value measured by the SDSS photometric pipeline for the r-band exponential disc fit. This reduces the number of free parameters in our model to seven. We down-weight the contribution of the velocity dispersion in our goodness-of-fit metric (equation 11) using $W = 0.2$. This is the same value Green et al. (2014) used in their analysis in order to achieve a high quality fit to the velocity field even if the selected velocity dispersion model can not describe the observed data very well.

In order to get a better estimate of the uncertainties on the model parameters we use a Monte Carlo method. For each galaxy, we compute the standard deviation of the best-fitting parameters over 100 simulated velocity and velocity dispersion fields. The simulated fields are created by perturbing the pixel values of the actual data based on the measurement error maps. In Fig. 4 we plot the residuals between our best fits and the best fits of Green et al. (2014).

Due to an error, the values of $r_{\rm t}$ reported in Green et al. (2014) do not account for the spatial sampling of the instrument which is 0.7 arcsec pixel^{-1} (A. Green, private communication, 2014), therefore a correction factor of 0.7 is applied prior the residual calculation. A small scatter is observed in the $r_{\rm t}$ residuals (Fig. 4A). This is because Green et al. (2014) used a fixed exponential surface brightness profile (i.e., constant scale radius of 7.5 pixels) for all the galaxies in the sample, while we used a varying one based on the scale radii found in the SDSS database (r-band photometry). However, the change to a varying scale radius does not affect the results and science conclusions of Green et al. (2014).

A very good agreement is observed in the plot of the circular velocity (Fig. 4B). For the circular velocity we used the velocity evaluated at a radius of 2.2 $\times$ scale lengths on the fitted rotation curve profile (Freeman 1970).

Our velocity dispersion results appear to be in a good agreement with the previous work (Fig. 4C). However, a few inconsistencies are observed. This is expected since the velocity dispersion in our case is derived from the fitting procedure, while the velocity dispersion in Green et al. (2014) is calculated using intensity weighted mean of the observed velocity dispersion map ($\sigma_{\rm m}$, Law et al. 2009). Furthermore, the intrinsic velocity dispersion of some galaxies is too small to be measured given the instrument's spectral resolution and signal-to-noise ratio, thus our method converges to unrealistically small values (Fig. 4C, empty circles). For a complete list of the fitting results see Table C2.

5 PERFORMANCE RESULTS

In this section we investigate the performance of our fitting method and assess the benefits of using the graphics processing unit. In order to obtain a better picture of the acceleration achieved by the GPU, a series of tests are performed on different hardware configurations using CPU-based and GPU-based versions of our code. We highlight good and bad uses of the GPU, we investigate how the performance scales

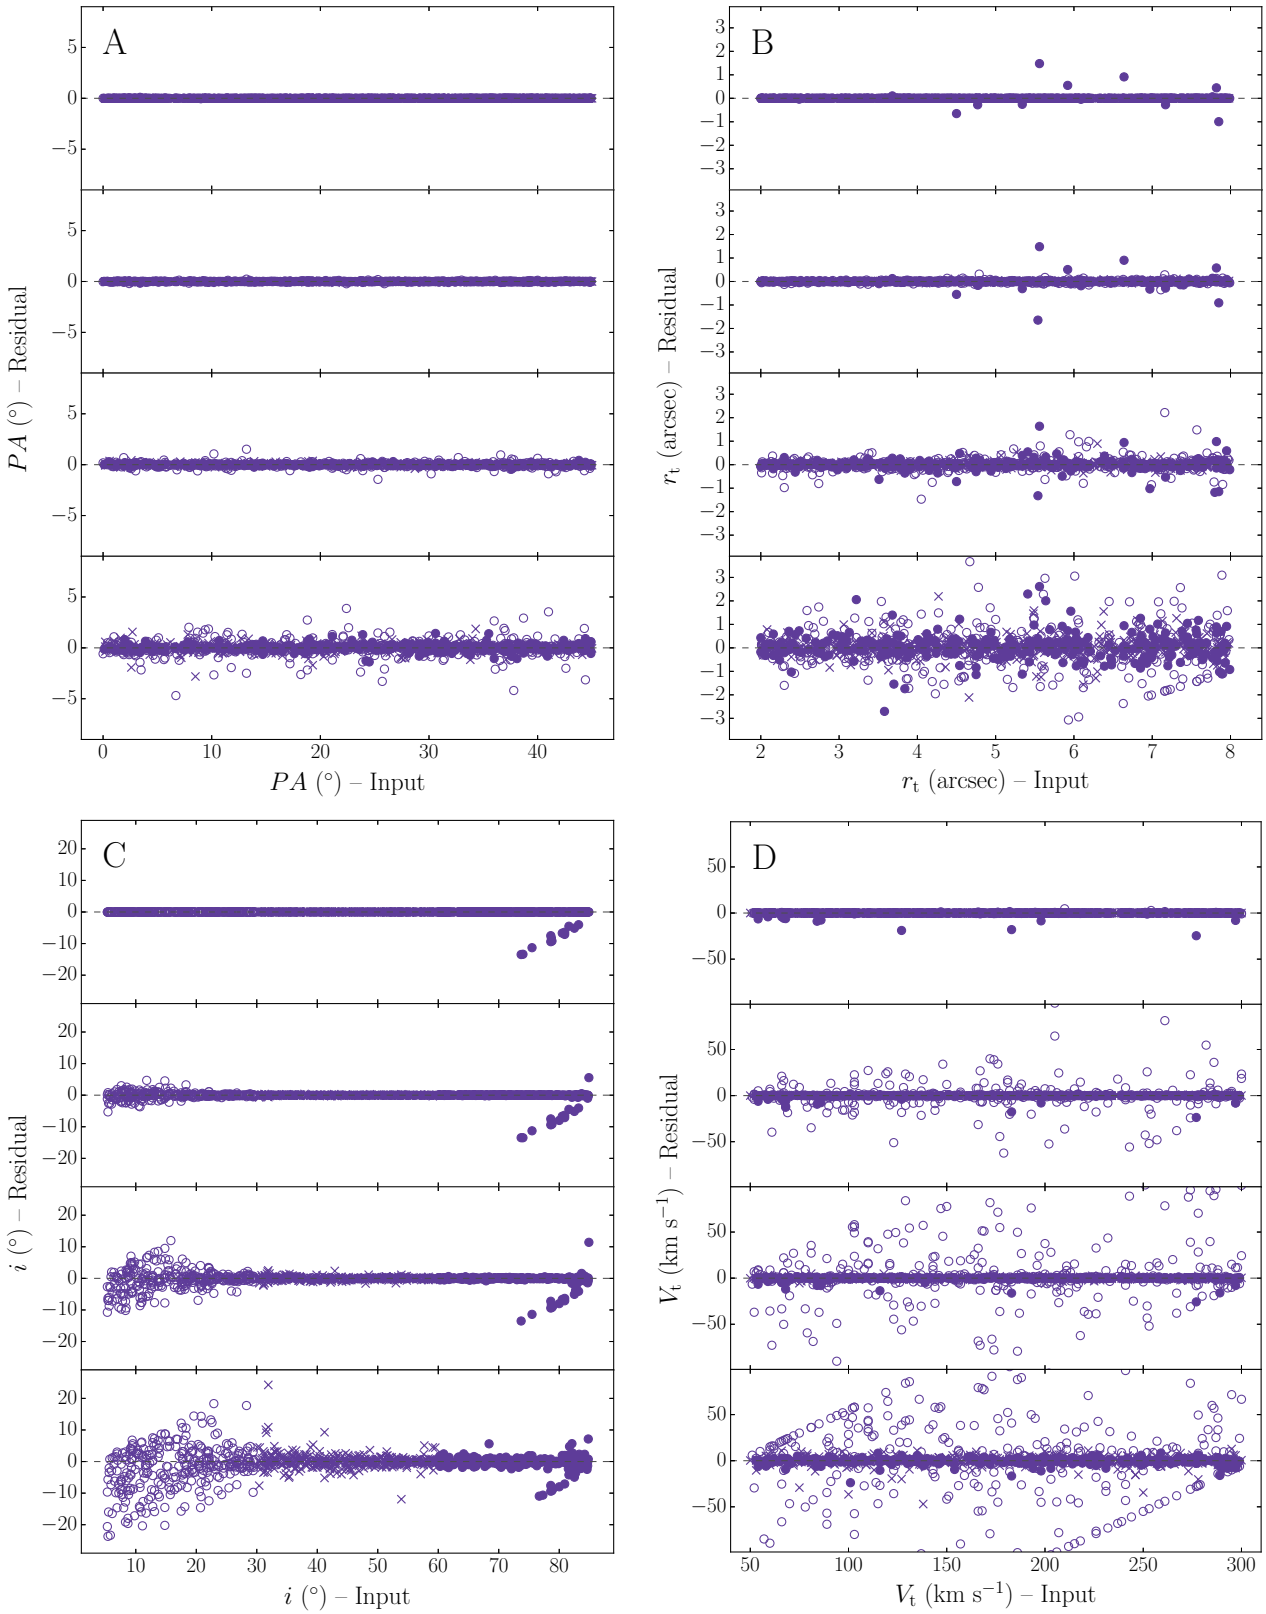

Figure 2. Fitting results of the mock data with different levels of noise using the Levenberg–Marquardt Algorithm. We plot the residuals of the fitted parameters as a function of the input parameters. (A) Position angle. (B) Turn-over radius. (C) Disc inclination. (D) Asymptotic circular velocity. In each of the plots, the rows from top to bottom correspond to datasets with 0, 1, 5, and 20 km s^{-1} noise on the velocity map, and 0, 1.4, 7, and 28 km s^{-1} noise on the velocity dispersion map. Different markers are used according to the input inclination: open circles for $5^\circ \leqslant i_{\rm Input} < 30^\circ$, crosses for $30^\circ \leqslant i_{\rm Input} < 60^\circ$, and filled circles for $60^\circ \leqslant i_{\rm Input} < 85^\circ$. The error bars are omitted for the sake of clarity.

Table 3. Number of fits with a parameter residual smaller than the specified threshold for the mock dataset for a total of 1000 zero-noise mock observations.

Optimization method	x_o (< 1 pixel)	y_o (< 1 pixel)	PA (< 1 °)	i (< 1 °)	r_t (< 1 pixel)	V_t (< 5 km s^{-1})	σ_{model} (< 1 km s^{-1})
LMA	1000	1000	1000	987	999	990	992
NS	1000	1000	1000	1000	1000	1000	1000
NG	1000	1000	998	1000	953	990	971

Table 4. Median of the residuals between actual and fitted parameters for the zero-noise mock dataset.

Optimization method	x_o (pixel)	y_o (pixel)	PA (°)	i (°)	r_t (pixel)	V_t (km s^{-1})	σ_{model} (km s^{-1})
LMA	0.7×10^{-3}	0.6×10^{-3}	0.1×10^{-2}	0.5×10^{-2}	1.3×10^{-3}	2.2×10^{-2}	1.6×10^{-3}
NS	0.7×10^{-3}	0.7×10^{-3}	0.1×10^{-2}	0.5×10^{-2}	1.4×10^{-3}	2.5×10^{-2}	1.8×10^{-3}
NG	–	–	0.9×10^{-1}	–	1.1×10^{-2}	9.8×10^{-2}	9.7×10^{-2}

Table 5. Number of fits with a parameter standardised residual within a specified threshold for a total of 1000 mock observations.

Optimization method	Noise (v/σ) (km s^{-1})	x_o ($1\sigma/2\sigma$)	y_o ($1\sigma/2\sigma$)	PA ($1\sigma/2\sigma$)	i ($1\sigma/2\sigma$)	r_t ($1\sigma/2\sigma$)	V_t ($1\sigma/2\sigma$)	σ_{model} ($1\sigma/2\sigma$)
LMA	1/1.4	657/939	663/943	709/957	671/934	676/937	660/927	684/944
NS	1/1.4	638/935	646/942	688/938	663/941	683/939	649/933	683/947
NG	1/1.4	1000/1000	1000/1000	991/998	1000/1000	794/896	766/878	988/993
LMA	5/7	694/948	661/949	676/948	684/925	707/941	662/910	698/943
NS	5/7	660/935	641/942	657/950	655/929	679/944	647/926	682/944
NG	5/7	1000/1000	1000/1000	935/967	1000/1000	697/877	636/812	912/955
LMA	20/28	657/951	689/959	698/961	654/879	651/930	632/849	702/945
NS	20/28	642/935	678/950	684/946	641/923	635/911	657/921	678/941
NG	20/28	1000/1000	1000/1000	867/950	1000/1000	708/933	724/912	803/951

with the size of our data and we comment on the memory requirements.

5.1 Hardware and software configuration

To understand how the performance of our method compares to CPU-based approaches, in addition to the main GPU implementation we have developed a single-threaded and a multi-threaded version of our code which runs on the CPU.

For the multi-threaded CPU implementation we used the OPENMP[4] library (Dagum & Menon 1998) to parallelize the code, and we chose a thread count higher than the number of available CPU cores in order to hide the memory access latency (Kurihara et al. 1991). The FFTW3[5] library (Frigo & Johnson 2005) was used to perform the FFTs on the CPU. All the allocations on the main system memory were done using a 16-byte alignment in order to give the compiler and the FFTW3 library the chance to use the SIMD capabilities of the CPU (Eichenberger et al. 2004). For the GPU implementation of our method we used the Nvidia Compute Unified Device Architecture (CUDA[6]) library to parallelize the code and the CUFFT[7] library to perform the FFTs.

The performance tests was run on the following four hardware configurations:

(1) 1 × Intel Xeon E5-2660 in single-threaded mode (ST)
(2) 2 × Intel Xeon E5-2660 in multi-threaded mode (MT)
(3) 1 × Nvidia Tesla C2070
(4) 1 × Nvidia Tesla K40c

Setups 1 and 2 are CPU-based with the first setup using only a single core while the second uses all the available cores from both CPUs. The Intel Xeon E5-2660 processor consists of 8 cores operating at 2.2GHz resulting in a theoretical peak processing power of ∼ 140 GFLOPS. The memory of the CPU-based setups was 64 gigabytes with a maximum theoretical bandwidth of 51.2 GB/s.

The last two setups, 3 and 4, are GPU-based and they represent the low- and high-performance ends of the GPU hardware spectrum respectively. The Tesla C2070 unit consists of 448 cores at 1.15GHz resulting in a theoretical peak processing power of ∼ 1TFLOPS, coupled with 6 gigabytes of memory of 144 GB/s bandwidth. The Tesla K40c unit consists of 2880 cores at 0.75GHz resulting in a theoretical peak processing power of ∼ 4.5TFLOPS coupled with

[4] http://www.openmp.org
[5] http://www.fftw.org
[6] http://developer.nvidia.com/about-cuda
[7] http://developer.nvidia.com/cufft

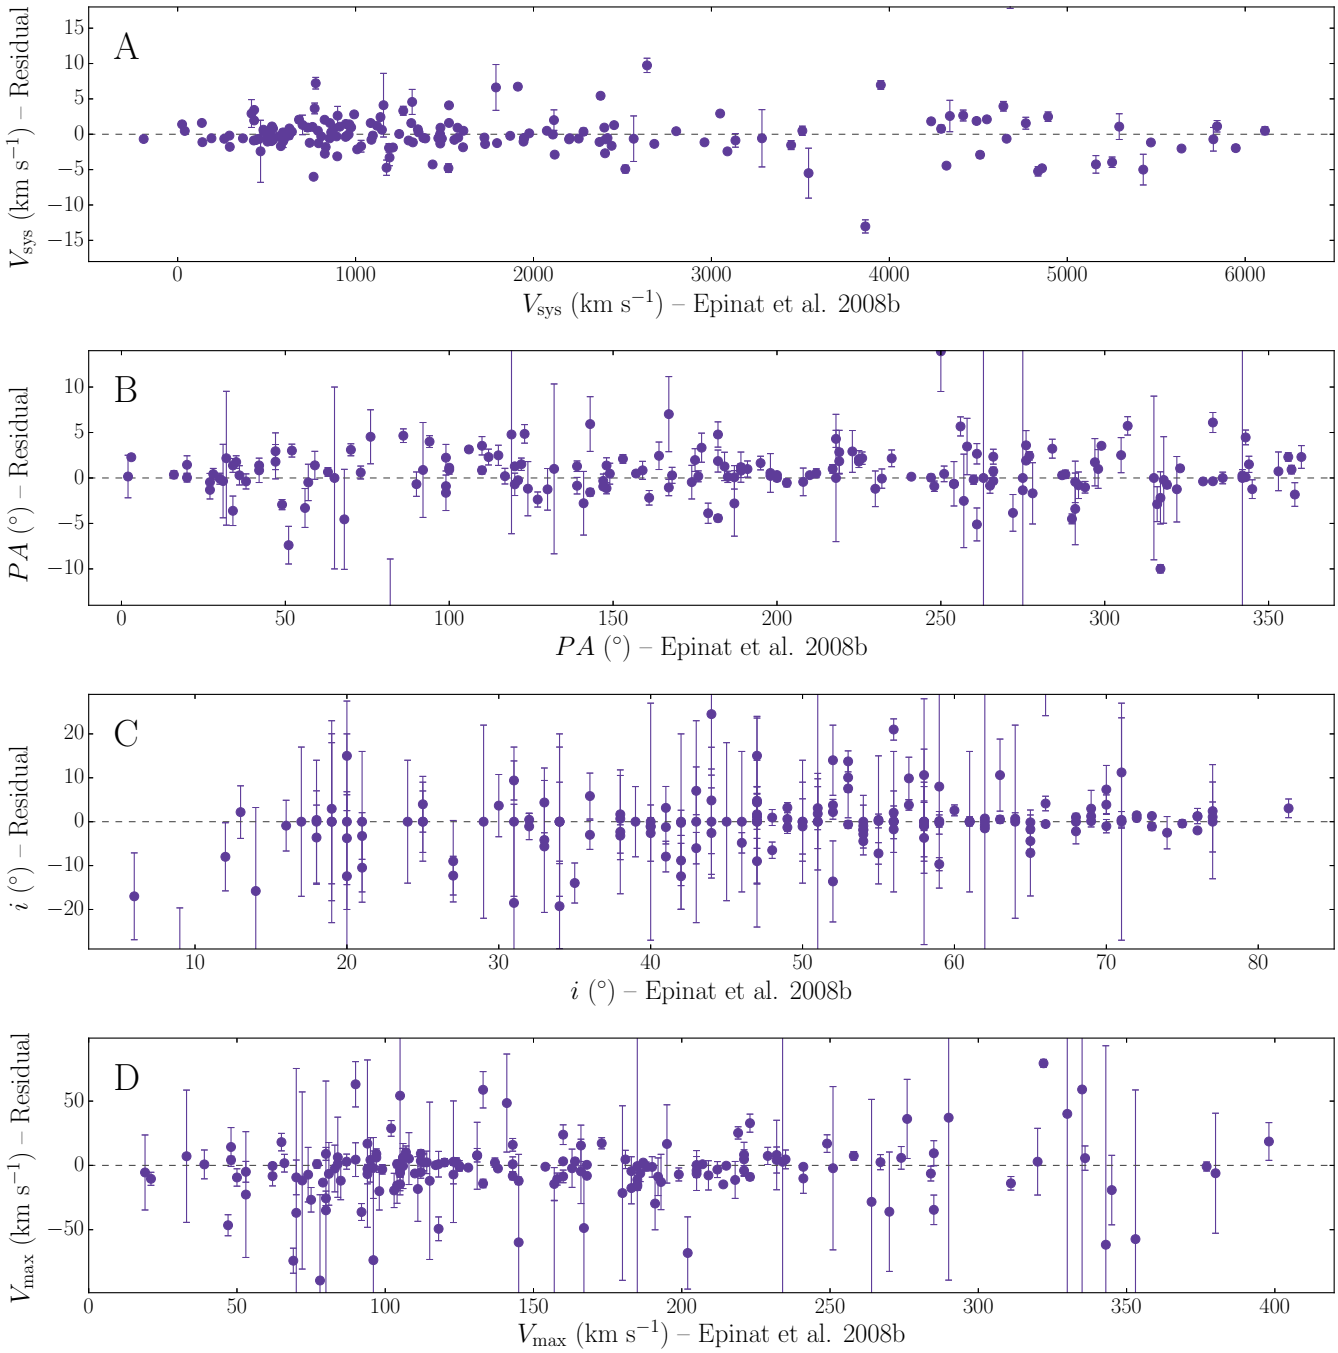

Figure 3. Fitting results of the GHASP sample. We plot the residuals of the fitted parameters as a function of the kinematic analysis results found in Epinat et al. (2008b). (A) Systemic velocity. (B) Position angle. (C) Disc inclination. (D) Maximum circular velocity.

12 gigabytes of memory of 288 GB/s bandwidth. The two Tesla units communicate with the CPU via the PCIe bus. Tesla C2070 uses a x16 PCIe version 2 slot with a theoretical bandwidth of 8 GB/s, while Tesla K40c uses a x16 PCIe version 3 slot with a theoretical bandwidth of 15.75 GB/s.

5.2 Performance analysis

To evaluate the performance of our method we measure the runtime of each of the steps involved in a fitting iteration. Since the number of operations in each step is heavily dependent on the size of the data, we create mock observations using a fixed set of morphological and kinematic parameter values but vary the pixel count. The analysis is performed for both the intensity cube and velocity field modelling strategies.

For the intensity cube modelling strategy we perform the analysis three times: one with a fixed pixel count (100 pixels) across the velocity dimension but varying the spatial size; one with a fixed spatial size (100 × 100 pixels) but varying the number of slices in the cube; and one with a

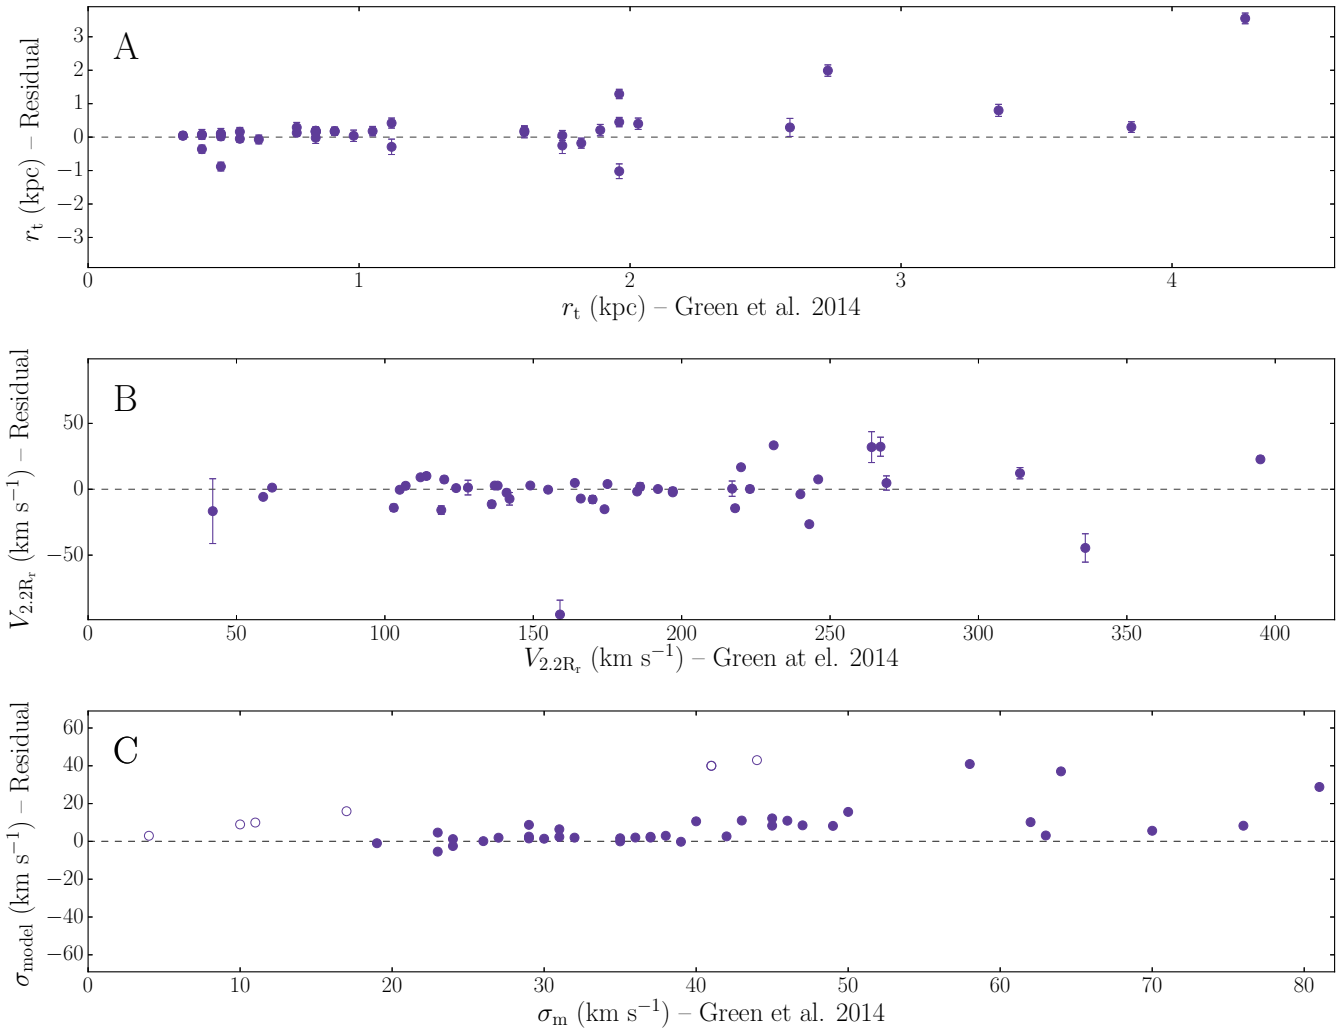

Figure 4. Fitting results of the DYNAMO sample. We plot the residuals of the fitted parameters as a function of the kinematic analysis results found in Green et al. (2014). (A) Turn-over radius. (B) Circular velocity at a radius of 2.2 × r-band exponential disc scale lengths. (C) Fitted velocity dispersion ($\sigma_{\rm model}$) against the velocity dispersion ($\sigma_{\rm m}$) found in Green et al. (2014). The open circles correspond to galaxies for which our fitting procedure converged to unrealistically small values.

fixed cube size ($50 \times 50 \times 50$ pixels), but varying the upsampling factor across all three dimensions. The size of the PSF kernel is set to $16 \times 16 \times 16$ pixels (for a seeing of ~ 2 kpc at a redshift of ~ 0.1). The above resolution is similar to the resolution of the DYNAMO observations.

The performance analysis results for the both modelling strategies are shown in Fig. 5, 6, 7, and 8. In most of the figures we observe the same trend: The total execution time of each step increases with the number of pixels in our data but also decreases as we move towards theoretically faster hardware. The GPU can be up to ~ 100 times faster than a single CPU core and up to ~ 10 times faster than a multi-threaded double CPU implementation. The intensity cube and velocity field model evaluation steps (Fig. 5A, 6A, and 8A) clearly demonstrate this substantial performance increase. Nevertheless, in some of the plots we observe performance behaviours that are worth comment.

For the FFT-based 3D convolution (Fig. 5B and 6B) we notice a step-like shape in the performance curves, which is expected since we always round-up the dimensions of the intensity cube model to the closest power of two. The FFT algorithms used by the CUFFT and FFTW3 libraries are highly optimized for input sizes that can be factored as $2^a \times 3^b \times 5^c \times 7^d$ (where a, b, c, and d are non-negative integers). In general the smaller the prime factor, the better the performance, i.e., powers of two are the fastest. Our method favours sizes of power of two, but for sizes more than 512, multiples of 256 are used. For example, an input size of 700 will be padded to 768 ($2^8 \times 3^1$). By doing so we conserve a significant amount of memory while also achieving good performance.

The intensity cube model downsampling step (Fig 7), although it has a very low arithmetic intensity and is dominated by memory transfers, it is still much faster on the GPU. This is expected since the GPU has a superior memory bandwidth compared to the CPU.

As a first attempt to extract the velocity and velocity dispersion maps from the modelled intensity cube, we fitted Gaussian profiles to the cube's spectra using the MPFIT library. Because MPFIT is a CPU-based library, the only opportunity for GPU parallelization is during the Gaussian

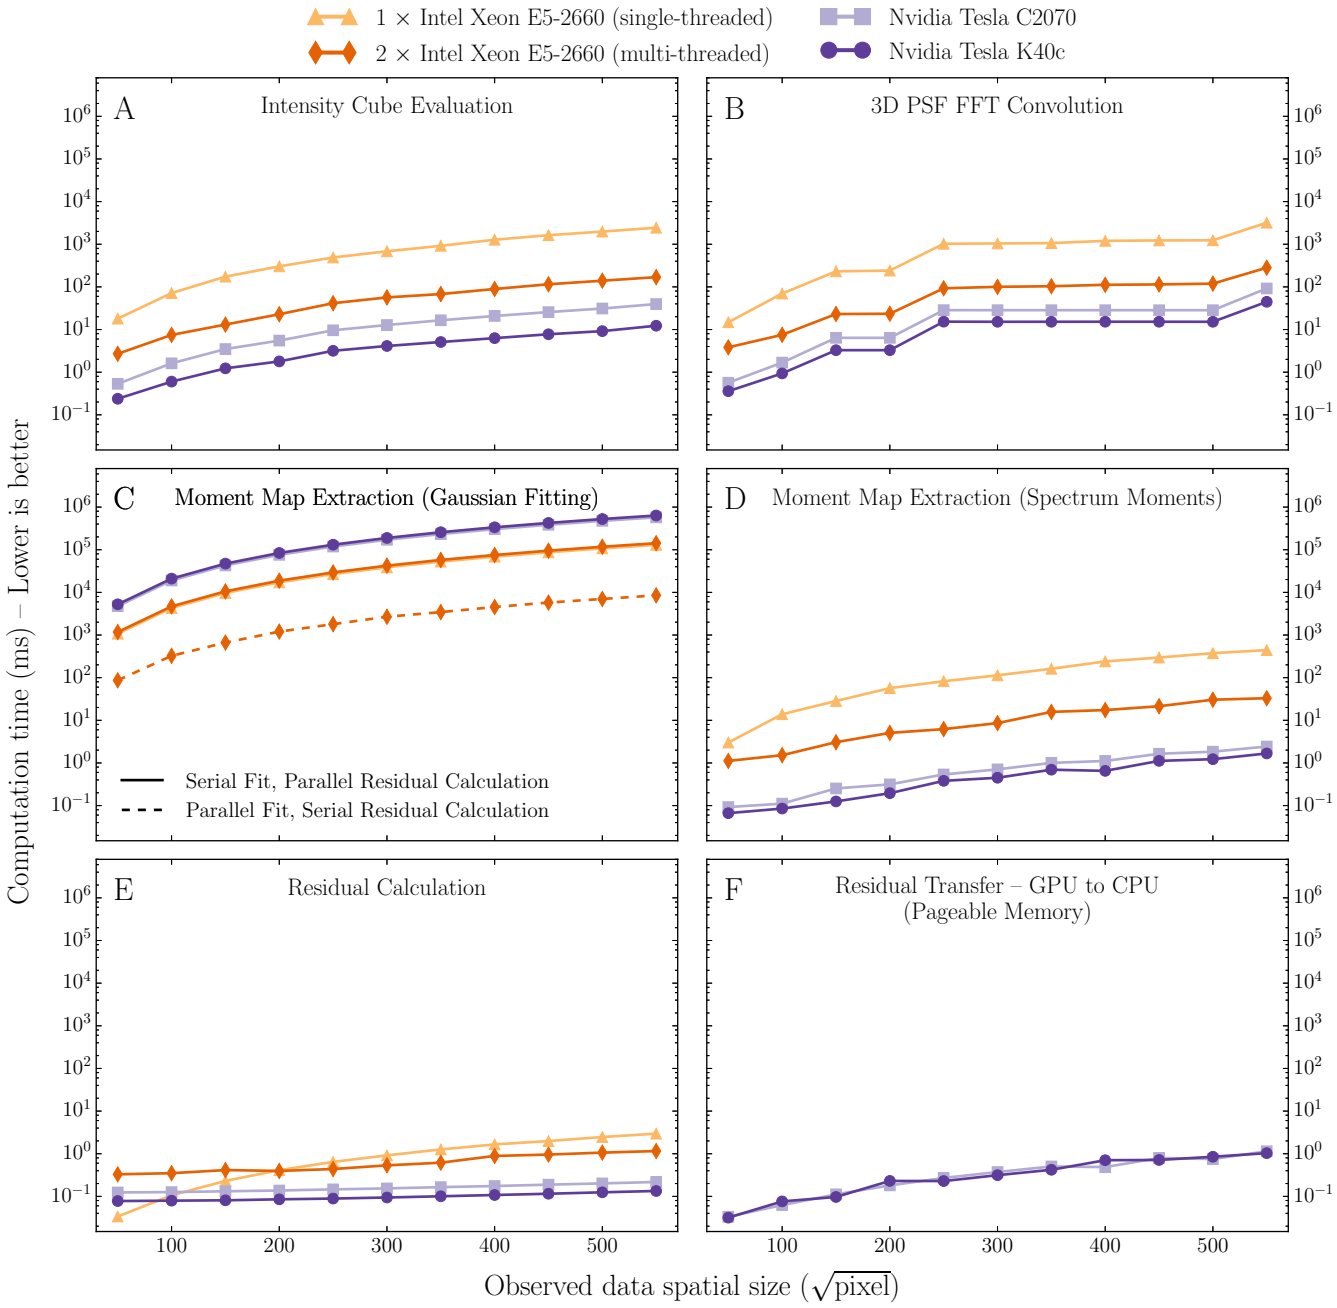

Figure 5. Performance comparison of four different hardware configurations for the steps involved in a single iteration of the fitting procedure using the intensity cube modelling strategy and a varying spatial pixel count. For this comparison we used no upsampling ($u = 1$) for all three dimensions of the intensity cube model, and a fixed number of 100 pixels (128 after padding) for the velocity dimension. (A) Intensity cube model evaluation. (B) FFT-based convolution with the 3D PSF kernel. (C) Velocity and velocity dispersion map extraction by fitting Gaussian profiles on the cube's spectral lines. (D) Velocity and velocity dispersion map extraction by calculating the first- and second-order moments of the cube's spectra. (E) Calculation of the residual velocity and velocity dispersion maps. (F) Residual map transfer from the GPU to the CPU memory.

profile evaluation and the calculation of the residuals. This approach has a few major performance issues. Since each spectral line of the model cube consists of only a few tens of pixels the majority of the GPU cores remain inactive and the GPU is underutilised. In addition, although the memory access is coalesced, the arithmetic intensity is so low that it makes the procedure memory bound. As a result of these issues, the GPU implementation of this step is slower than the CPU version (Fig. 5C and 6C). Furthermore, an extra overhead is created because the LMA logic of MPFIT runs on the CPU serially. This overhead overshadows the time requirements of the parallelized residual calculation and as a result we observe a negligible performance difference between the high- and low-end GPUs and between the single- and multi-

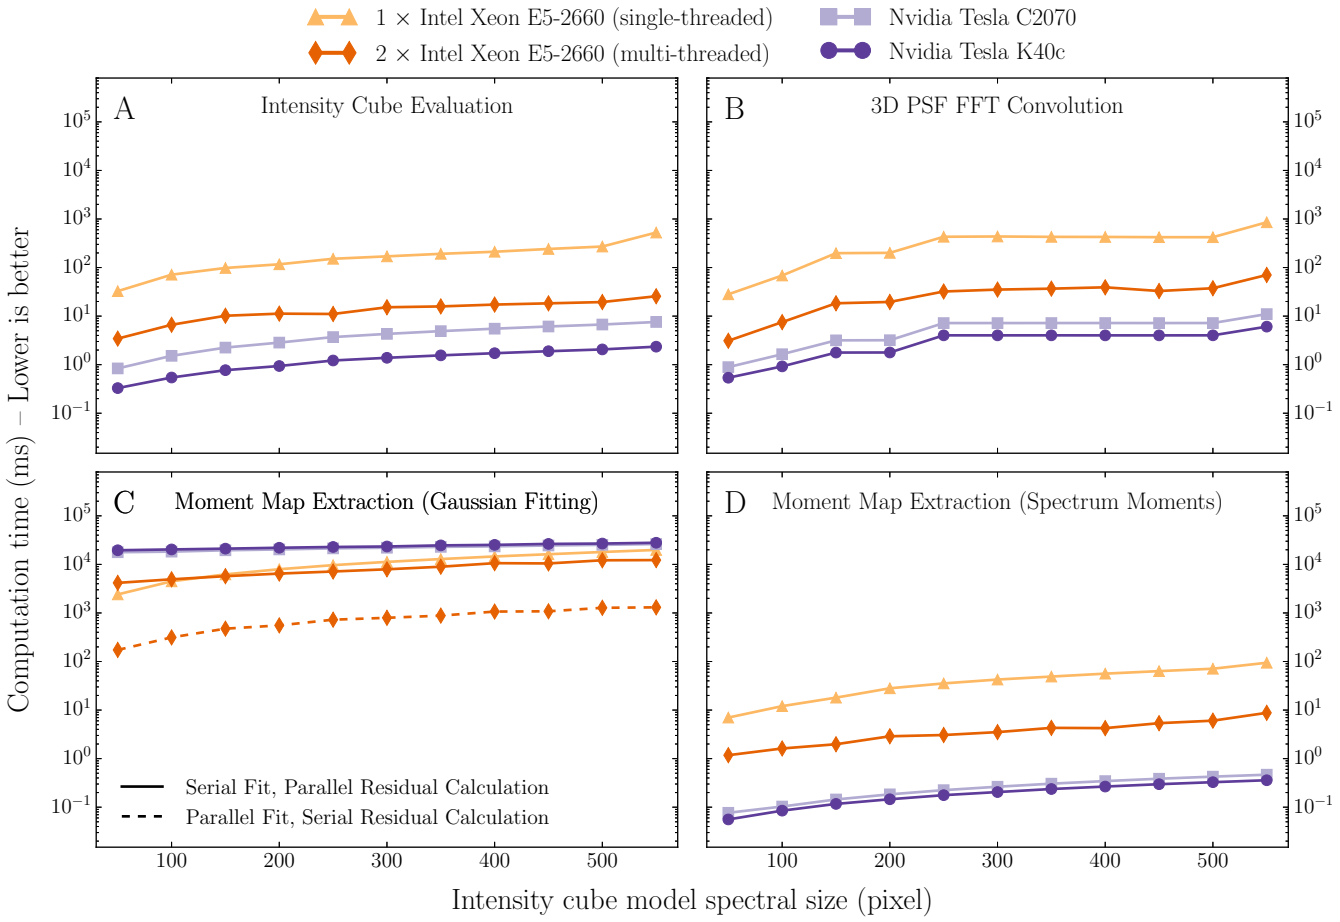

Figure 6. Performance comparison of four different hardware configurations for all the steps involved in a single iteration of the fitting procedure using the intensity cube modelling strategy and a varying spectral pixel count. For this comparison we used no upsampling ($u = 1$) for all three dimensions of the intensity cube model, and a fixed number of 100×100 pixels (128×128 after padding) for the spatial dimensions. (A) Intensity cube model evaluation. (B) FFT-based convolution with the 3D PSF kernel. (C) Velocity and velocity dispersion map extraction by fitting Gaussian profiles on cube's spectral lines.

threaded CPU configurations. A much better performance for this step was achieved by performing the Gaussian profile evaluation and residual calculation serially but fitting multiple spectral lines in parallel on the CPU. To do the same on the GPU would require a GPU-specific implementation of the Levenberg–Marquardt algorithm. Zhu & Zhang (2013) showed that implementing the LMA to run solely on the GPU can achieve a performance significantly superior to the CPU but such code is not yet publicly available.

Finally, we decided to approach the problem of the velocity and velocity map extraction differently by calculating the first- and second-order moments of the cube's spectral lines. This method is much simpler compared to an iterative one (e.g., the Levenberg–Marquardt algorithm) and also highly parallelizable. Consequently, a parallelized implementation of this approach on the GPU was trivial and we were able to achieve a substantial performance boost compared to the CPU implementation (Fig. 5D and 6D).

The residual calculation step features an atypical performance behavior. The computational time remains almost constant regardless of the number of pixels in our dataset (Fig. 5E and 8B). During the development of our GPU code we wrote a series of reusable CUDA programs (*kernels*). Instead of writing a kernel which calculates the residual of the velocity and dispersion maps specifically, we wrote a kernel that calculates the residual of any image and we call it twice, once for each map. Each time we call a CUDA kernel there is an associated overhead. If the computation time requirements of our kernel are very low then they are overshadowed by that overhead. In our case, not only are the computational requirements of the kernel low, but also the kernel call overhead occurs twice. However, the computational time of this step is negligible compared to the computational time requirements of the rest of the steps in the iteration.

The runtime of the residual map transfer from the GPU to the CPU memory is much higher than the expected (Fig. 5F). It also appears to be the same for both GPU setups, even though the Tesla K40c system has twice the theoretical bandwidth of the Tesla C2070 system. While this poor performance is not an important issue for the intensity cube modelling strategy, it becomes noticeable when the velocity field modelling strategy is used (Fig. 8C). We were able improve the runtime of the residual map transfer by

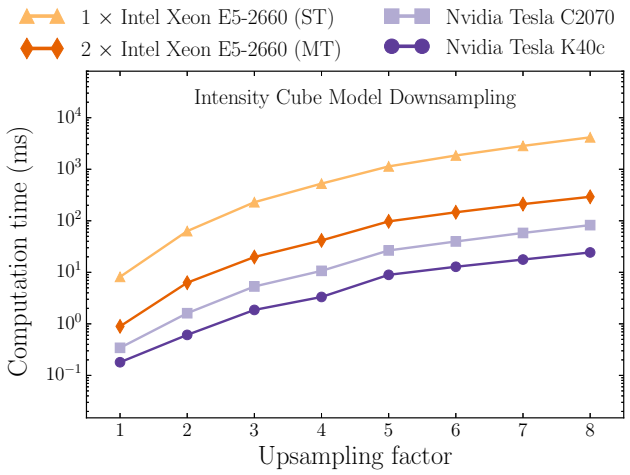

Figure 7. Performance comparison of four different hardware configurations for the downsampling step involved in a single iteration of the fitting procedure using the intensity cube modelling strategy and a varying upsampling factor (u). For this comparison we vary the upsampling factor simultaneously across all three timensions, while the initial size of the intensity cube model is fixed to $50 \times 50 \times 50$ pixels.

making use of the CUDA pinned memory (Fig. 8D). By using CUDA pinned memory we avoid an extra copy between the CPU's virtual and physical memory that the CUDA driver has to make in order to ensure that all CPU↔GPU transfers are done through CPU's non-pageable memory (Fatica 2009).

To get a better idea about the time required to fit real observations, we fit the GHASP and DYNAMO data using different optimization techniques and hardware configurations. An upsampling factor of $u = 2$ across all three dimensions was used while modelling the DYNAMO data. In Fig. 9 we plot the total time spent to fit all the galaxies from both samples. For the DYNAMO sample, the speedup achieved by the GPU is between ~ 10 and ~ 100 depending on the hardware configuration. When fitting the GHASP sample, however, the speedup difference between the two GPUs is not as big as expected. This happens because the runtime requirements of the velocity field modelling strategy are comparable to the runtime requirements of the optimizer. In Table 6 we quote the average number of iterations and the time required to fit a galaxy, but we also report how much of that time is spent on the optimization logic.

In general, fitting real observations tends to be slower than fitting artificial observations because most of the times the optimizer requires more iterations before converging to a solution. This can be explained due to the fact that, unlike real data, the simulated observations used in this work do not include any random motions or correlated noise, and thus the resulting posterior distribution of the parameters is usually well behaved.

5.3 Memory considerations

Apart from our main goal to accelerate the fitting procedure, it is also desirable to keep the memory requirements of our method at a low level. When running CPU-based software this is hardly an issue because of the large main system memory available on today's computers. In this work we did not encounter any GPU memory problems because we used enterprise-grade Tesla GPU units which feature a relatively large amount of memory (up to 12GB). However, it is desirable to run our method on the commodity versions of these cards, which have a much more limited amount of memory (e.g., 2GB).

During the runtime of our code most of the utilised memory is occupied by:

(1) The data of the intensity cube model.
(2) The data of the intensity cube model in Fourier space.
(3) The data of the PSF kernel in Fourier space.
(4) The data buffers allocated by the CUFFT library.

The expected size (in bytes) of a 3D dataset such as the intensity cube model is:

$$M_{3D} = N_z \times N_y \times N_x \times 4, \tag{12}$$

where N_z, N_y, N_x the length of each dimension in pixels and the size of a single precision real number is 4 bytes. The expected size (in bytes) for the same 3D dataset but in frequency space is calculated using:

$$M_{FFT,3D} = N_z \times N_y \times (N_x/2 + 1) \times 8, \tag{13}$$

for single precision complex numbers.

CUFFT allows us to perform the FFT in-place without allocating additional space for the frequency space data. In order to do so we have to make sure that the memory space we allocate for the data is padded to the size required to accommodate the data in frequency space.

To perform an FFT, the CUFFT library has to create a so-called *plan*. This object contains all the data CUFFT needs to compute the FFT. Internally the plan allocates memory buffers which are used for interim results. The size of these buffers depends on the size of the FFT and it can be quite large. Kauker et al. (2010) showed that CUFFT needs about four times the space of the input for its intermediate results for real-to-complex transforms when transforming two-dimensional data. The same plan can be applied to different datasets as long as they have the same length.

To get a better picture of the memory requirements of the CUFFT library we perform a series of three-dimensional FFTs followed by inverse FFTs. More specifically, we generate 3D datacubes of different sizes and for each one of them we perform an in-place real-to-complex 3D FFT followed by a complex-to-real inverse 3D FFT. Since two different types of transforms are performed for each dataset, two plans are needed. Both plans appear to occupy the same amount of memory. The results are plotted in Fig. 10.

For datasets with a number of pixels below $\sim 200^3$ the memory requirements of the two FFTs are much smaller than the memory available on todays low-range commodity GPUs. At a resolution of $\sim 550^3$ the required size exceeds the typical memory size available on a mid-range commodity GPU. In addition, we notice that the plans require more memory than the actual data itself. The maximum FFT size the CUFFT library can compute is ~ 512 million single precision elements, which corresponds to a 3D FFT of $\sim 800^3$ or a plan of ~ 2 gigabytes.

Here we present an easy way to reduce the size of the

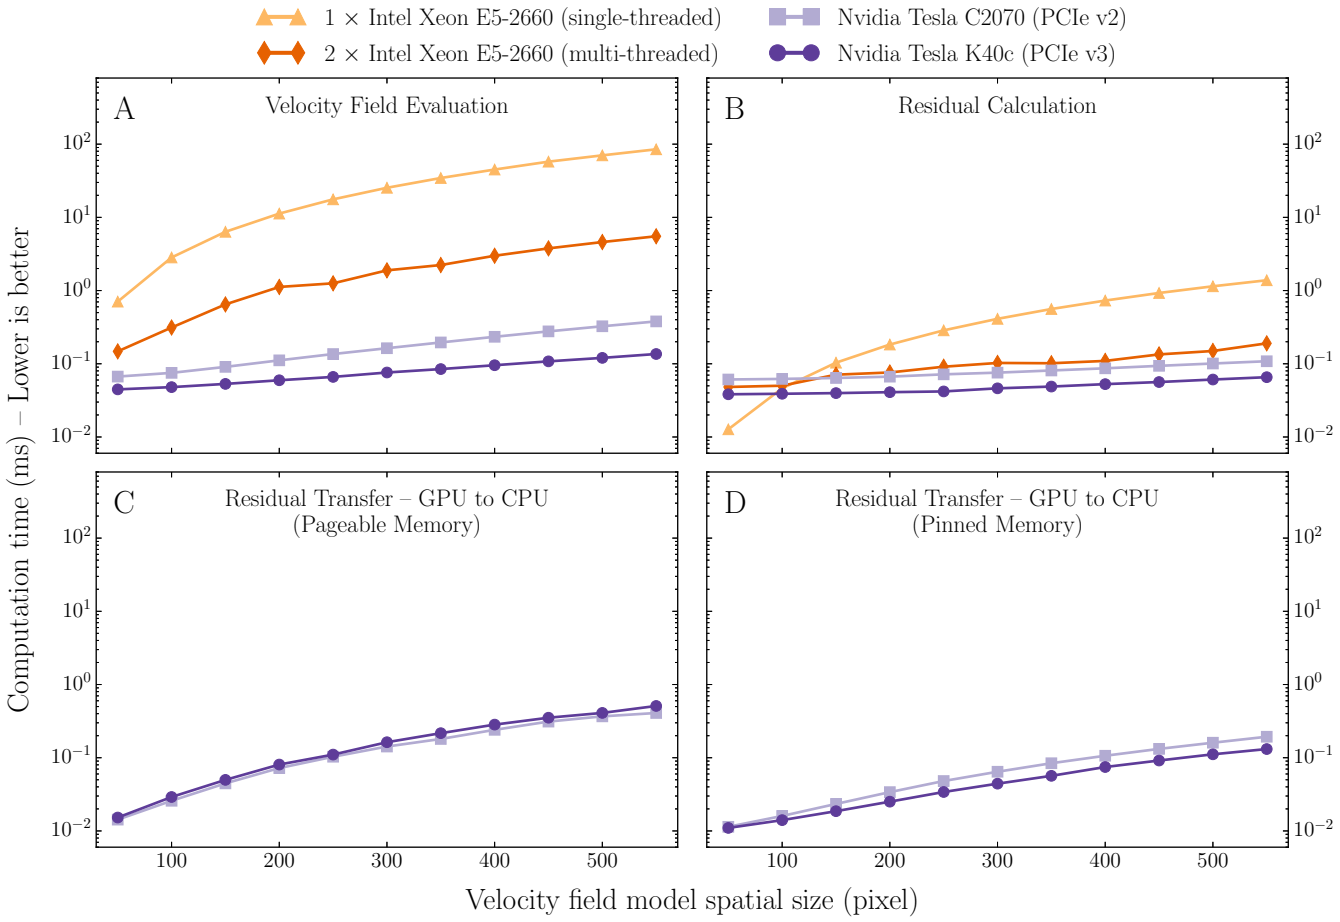

Figure 8. Performance comparison of four different hardware configurations for all the steps involved in a single iteration of the fitting procedure using the velocity field modelling strategy. (A) Velocity field model evaluation. (B) Calculation of the residual velocity map. (C) Residual map transfer from the GPU to the CPU memory. (D) Residual map transfer from the GPU to the CPU memory using CUDA pinned memory.

FFT plans and enable us to perform large 3D FFTs on commodity GPUs. The 3D FFT is decomposed into a series of 2D FFTs across the two dimensions of the data followed by a series of 1D FFTs across the third dimension. Although this results in much smaller memory requirements for the plans, the total computation time increases. This is expected since we now have to call the appropriate CUFFT functions for plan execution multiple times instead of one, which creates an additional overhead. We overcome this issue by using CUFFT's batched plans which can perform multiple FFTs of the same length with one function call. However, the size of the batched plans increases with the number of transforms in the batch. Hence, in order to achieve balance between performance and memory requirements, one has to carefully select the batch size of the plans.

Fig. 11 shows how the performance and the plan size vary when calculating 512 two-dimensional FFTs with different batch configurations. A compromise of $10 - 20$ per cent in performance can decrease the memory requirements by one order of magnitude. Finally, the 1D FFT that has to be applied across the third dimension of the data in order to calculate the 3D FFT, behaves in a similar manner hence it is not discussed here.

To lower the memory requirements even further, we perform the convolution in the frequency domain using the point (2D) and line (1D) spread functions instead of the 3D PSF kernel. As a result, the memory requirements for the PSF data are reduced from $M_{\rm FFT,3D}$ (equation 13) to:

$$M_{\rm FFT,2D+1D} = [N_{\rm y} \times (N_{\rm x}/2 + 1) + (N_{\rm z}/2 + 1)] \times 8. \quad (14)$$

To get a better understanding of the memory savings after applying the changes discussed in this section, we run the fitting procedure for low- and high-resolution mock data on a commodity laptop equipped with a 2GB Nvidia GeForce 750M GPU. The low-resolution data have the same resolution as the DYNAMO observations, while the high-resolution data have 700×700 spatial pixels and a velocity resolution of 15 km s^{-1}. For both datasets the upsampling factor was set to $u = 1$ for all three dimensions. The GPU memory requirements are now $\sim$ 2 megabytes down from $\sim$ 11 megabytes for the low-resolution dataset, and $\sim$ 0.7 gigabytes down from $\sim$ 3 gigabytes for the high-resolution dateset.

Table 6. Average number of iterations and time spent to fit a GHASP or DYNAMO galaxy using different optimization techniques and hardware configurations.

Survey	Optimizer	Hardware	$N_{\rm median}$	$N_{\rm mad}$	$t_{\rm opt}$ (seconds)	$t_{\rm total}$ (seconds)	$t_{\rm opt}/t_{\rm total}$
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GHASP	LMA	1 × Intel Xeon E5-2660 (ST)	234	63	0.77	7.48	10.35%
		2 × Intel Xeon E5-2660 (MT)	234	63	0.74	1.18	62.84%
		Nvidia Tesla C2070	242	70	0.83	0.89	93.56%
		Nvidia Tesla K40c	242	70	0.79	0.82	95.74%
	NS	1 × Intel Xeon E5-2660 (ST)	12030	4739	2.89	351.21	0.82%
		2 × Intel Xeon E5-2660 (MT)	11852	4832	3.01	24.96	12.07%
		Nvidia Tesla C2070	12207	4777	2.61	7.09	36.84%
		Nvidia Tesla K40c	13496	5818	3.20	5.16	62.09%
	NG	1 × Intel Xeon E5-2660 (ST)	232925	21175	31.12	6063.23	0.51%
		2 × Intel Xeon E5-2660 (MT)	232925	21175	32.70	412.87	7.92%
		Nvidia Tesla C2070	232925	21175	35.52	144.93	24.51%
		Nvidia Tesla K40c	232925	21175	30.48	58.80	51.83%
DYNAMO	LMA	1 × Intel Xeon E5-2660 (ST)	148	25	0.01	75.61	0.01%
		2 × Intel Xeon E5-2660 (MT)	148	25	0.01	11.73	0.08%
		Nvidia Tesla C2070	153	30	0.01	2.22	0.35%
		Nvidia Tesla K40c	153	30	0.01	0.69	1.11%
	NS	1 × Intel Xeon E5-2660 (ST)	10706	2230	1.18	5505.42	0.02%
		2 × Intel Xeon E5-2660 (MT)	9623	2869	1.30	797.38	0.16%
		Nvidia Tesla C2070	12308	3600	0.78	110.27	0.71%
		Nvidia Tesla K40c	10921	2423	0.90	51.02	1.76%
	NG	1 × Intel Xeon E5-2660 (ST)	42350	0	0.30	20373.03	0.00%
		2 × Intel Xeon E5-2660 (MT)	42350	0	0.45	2948.69	0.02%
		Nvidia Tesla C2070	38115	4235	0.18	699.01	0.03%
		Nvidia Tesla K40c	38115	4235	0.20	189.17	0.11%

(1, 2, 3) The dataset, optimization technique, and hardware configuration used in the fitting procedure. (4, 5) Median and median absolute deviation of the number of iterations performed. (6) Median of the time spent on optimization logic. (7) Median of the total runtime.

6 DISCUSSION

In the previous sections we demonstrated how the graphics processing unit can be used as an accelerator for the kinematic fitting procedure. Our method adopts a similar pattern to the one followed by most existing kinematic fitting software. Two different modelling strategies were explored: One uses a velocity field and it is more appropriate for high resolution observations, while the other one uses an intensity cube model and the knowledge of the point and line spread functions in order to account for beam smearing effects.

We focused on parallelizing the steps involved in the fitting procedure but did not attempt to parallelize the optimizer's logic on the GPU. This allowed us to use pre-existing optimization libraries and accelerated the development of the software implementation. Due to the image nature of the data involved, we were able to use data-parallel algorithms and exploit the power of the GPU's massively parallel architecture.

To gain a better idea of the acceleration the GPU has to offer we also implemented and ran our method on the CPU. In most cases, our tests showed a remarkable performance increase when the GPU was used. We achieved a maximum of ~ 100 times increase over a single-threaded CPU implementation and a peak of ~ 10 times increase over a multi-threaded dual CPU implementation.

Nevertheless, we encountered a case were the GPU performed worse than the CPU. This was the step where Gaussian profiles were fitted to each spectrum of the intensity cube model in order to extract the velocity and velocity dispersion maps. To achieve a good performance for this task we would need to implement an LMA optimizer that runs entirely on the GPU (Zhu & Zhang 2013), the development of which is still ongoing. Instead, to improve the performance of this step we followed a different approach using the cube's spectral line first- and second-order moments.

A peculiar performance behaviour was also noticed in the velocity field modelling strategy: Although the speedup achieved by the GPU for the velocity field evaluation step is remarkable, the same amount of speedup was not observed on the total time required to fit a GHASP galaxy. The reason for this behaviour is that the parallelized steps involved in the velocity field modelling strategy have very low runtime requirements and comparable to the time spent on the optimization logic.

By looking at the theoretical computational performance of our hardware (Section 5.1) we notice that in several cases we are below the expected speedup. This is not a surprise since the theoretical performance of the graphics processing unit is rarely achieved in real-world applications. In addition, achieving the peak performance of the GPU can require a high level of expertise and result in complicated or hard-to-read code (Ryoo et al. 2008). Nonetheless, the speedup achieved by our GPU-based method is remarkable.

Memory requirements are also an important factor when developing GPU software. Although the amount of

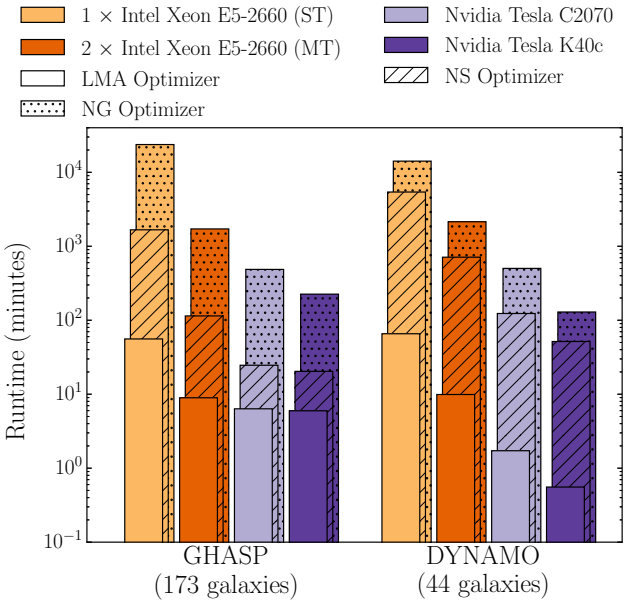

Figure 9. Total time spent to fit the entire GHASP and DYNAMO surveys using different hardware configurations and optimization techniques.

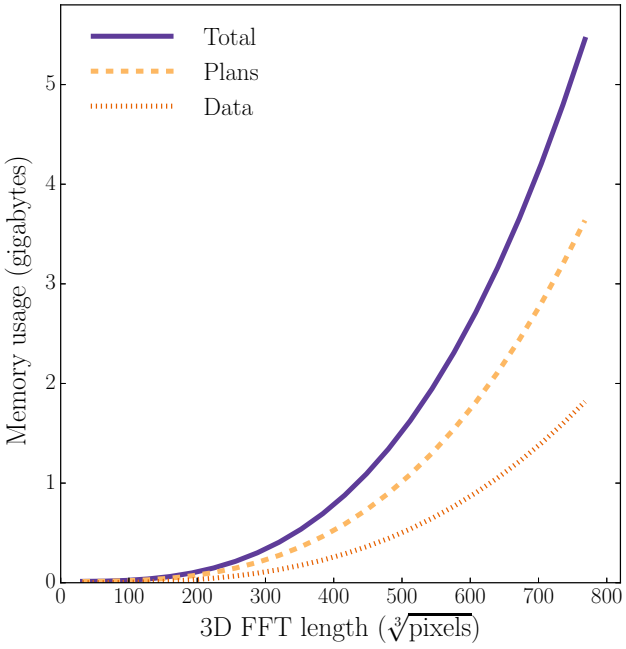

Figure 10. Total GPU memory consumption for varying-size three-dimensional real-to-complex FFT transforms followed by inverse FFT transforms using the CUFFT library. All transforms use single precision data and are done in-place.

memory on the GPU increases every year, so does the resolution of the data produced by the new generation astronomical instruments. We found that the FFT-based 3D PSF convolution step in the intensity cude modelling strategy can require a large amount of memory when working with high resolution data, and we proposed a simple solution to sub-

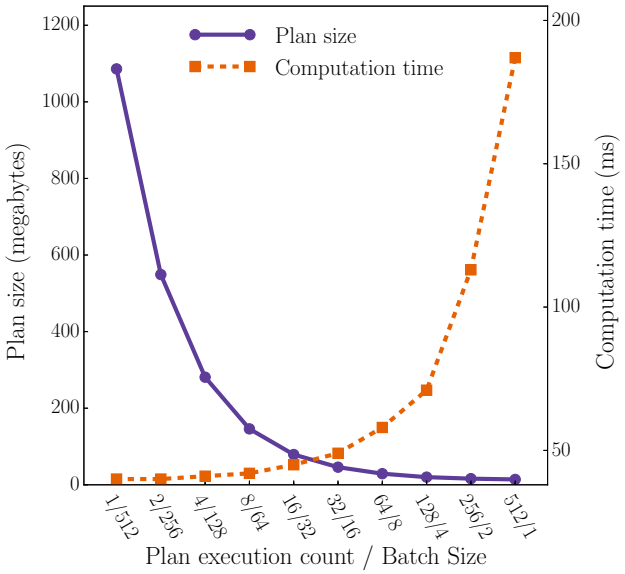

Figure 11. FFT plan memory consumption and execution time for two-dimensional 512^2 real-to-complex FFT transforms followed by inverse FFT transforms using the CUFFT library. All transforms use single precision data and are done in-place. A total of 512 2D FFTs are computed, each time with a different plan configuration.

stantially decrease the memory requirements using 2D+1D FFT and convolution decomposition (Section 5.3).

An issue one might face when developing software for the GPU is that existing CPU-codes has to be rewritten (Fluke et al. 2011). Although this was a bigger issue in the past, the situation has improved. Libraries and tools that ease the GPU code development or the porting of CPU code to the GPU have been developed. CUDA unified memory relieves the developer from managing the GPU memory (Landaverde et al. 2014). The THRUST[8] library (Kaczmarski & Rzazewski 2013) can be used to perform data parallel operations which run on the GPU using an interface that resembles parts of the C++ standard library. With OPENACC[9] one can port existing CPU code to the GPU by just adding a series of compiler directives (Wienke et al. 2012) to the code. We expect that technologies and tools like these will continue to improve and further decrease the effort required to write GPU software.

Another interesting point one has to consider is that the total cost of a GPU solution is much lower than the CPU cost for the same amount of computational power. Although in this work we used enterprise-level Tesla GPUs, the commodity versions of these cards, which have very similar performance capabilities and energy requirements, cost one order of magnitude less.

Our method was tested by fitting artificial observations with zero noise. We were able to recover the actual kinematic parameter values for $\gtrsim 98$ per cent of the galaxies with high accuracy (Table 3). In order to successfully recover the rest ~ 2 per cent the fitting procedure was re-run using a different optimizer configuration. This translates to better initial

[8] http://developer.nvidia.com/thrust
[9] http://www.openacc.org

parameter values or prior constraints for the LMA optimizer, lower tolerance threshold (*tol*) and more active points (N) for the NS optimizer, and higher grid resolution for the NG optimizer.

To fit the mock data using the NG optimizer we had to fix parameters $V_{\rm sys}$, $x_{\rm o}$, $y_{\rm o}$ and i to their true values in order to reduce the number of iterations required by the fitting procedure. This is still relevant of a real-world scenario since it is sometimes preferred to derive the kinematic centre and inclination of the galaxy from high-resolution photometry data rather than the velocity field (Epinat et al. 2010). By using the GPU we were able to fit the data in a reasonable amount of time regardless of the inefficiency of the algorithm.

Our method was also able to fit mock observations with low and high levels of noise with great success: Regardless of the noise level, ~ 68 and ~ 95 per cent of the the galaxy parameters were enclosed within the 1 and 2σ confidence intervals of our fits respectively.

Furthermore, our method verified successfully the kinematic modelling results of the low-redshift surveys GHASP and DYNAMO as found in the literature. Although in this work we only fitted data from optical surveys, our method can potentially be used at other wavelengths as long as an appropriate model is employed. For example, when fitting models to 21-cm data it is preferable to use a model that accounts for the warping observed at the outer parts of the neutral hydrogen layer in the majority of galaxies (Bosma 1981; Briggs 1990; García-Ruiz et al. 2002).

6.1 The degeneracy between $V_{\rm t}$ and i

Caution is advised when fitting near face-on galaxies, especially in the case of low-resolution data and/or observations with low signal-to-noise ratio. In such case, recovering the disc inclination from the kinematics can be quite problematic due to the near-degeneracy between the circular velocity and inclination.

To better demonstrate this degeneracy we generated four mock galaxies with varying inclinations (20°, 40°, 60°, and 70°) and constant circular velocity $V_{\rm t} = 200$ km s^{-1}. Gaussian noise of 5 km s^{-1} was added to the mock velocity and velocity dispersion maps. In Fig. 12 we plot the reduced χ^2 space as a function of the disc inclination (i) and circular velocity ($V_{\rm t}$). This figure illustrates that the degeneracy between the two parameters is quite strong for inclinations less then 40°.

To investigate whether the scatter observed in Fig. 2B and 2C is due to this degeneracy we plot the residual between the actual and the fitted product $V_{\rm t} \sin i$ (Fig 13). Indeed the combination of the two parameters reduces the scatter and their product is recovered with a good accuracy.

6.2 Moments vs Gaussian fitting

There are cases where the input velocity and velocity dispersion maps are extracted using the Gaussian fitting method but the modelling of the kinematics is done using the spectral moments approach due to the much lower runtime requirements and the simplicity of the implementation (e.g., Green et al. 2014; Section 4.3, this paper). To get a better picture on how the two methods compare to each other we create several mock observations with the same properties but varying the seeing size. Then we extract the velocity and velocity dispersion maps using both methods and we calculate the difference between them (Fig 14).

Under good seeing conditions we observe no difference between the two methods. However, as the size of the PSF increases, low-amplitude features appear on both maps. More specifically, for the maps extracted with the Gaussian fitting approach, the line-of-sight velocity gradient is steeper in the inner regions of the galaxy, while the velocity dispersion at the centre is slightly higher. Instead of increasing the PSF size, the same behaviour would be observed by reducing the turn-over radius of the galaxy.

To understand how the above differences can affect the modelling results we fit a model that uses the spectral moments approach to data that have been extracted using Gaussian fitting. The procedure is performed twice, with the data truncated at $r_{\rm trunc} = 1.0 \times r_0$ and $r_{\rm trunc} = 2.2 \times r_0$. The scale radius of the modelled galaxy has chosen to be $r_0 = 3.0 \times r_{\rm t}$ In Fig. 15 we plot the absolute error between the actual and the fitted parameters as a function of PSF FWHM.

We observe that when the flat part of the rotation curve is detected in the data (Fig. 15, left column), the error in the inclination, circular velocity, and intrinsic velocity dispersion is very small, even under bad seeing conditions. However, if the data are truncated at a much smaller radius (Fig. 15, right column), the error on the fitted parameters is noticeable, especially for observations with severe beam smearing. The errors on the rest of the model parameters are not presented here because they are negligible.

While the above suggest caution in selecting the right modelling approach for a given dataset, they are not necessarily an issue if the affected parameters are not essential for any resulting analysis. For example, if we are interested in the systemic velocity, kinematic centre or position angle of a galaxy, choosing one approach over the other will not affect our results. However, if we want to recover the disc inclination from a heavily beam-smeared observation truncated at small radius, we will have to choose our modelling approach carefully. In the case of the DYNAMO sample, since the galaxies are measured up to a large radius, both approaches yielded similar results.

Finally, we remind the reader that, as shown in Section 4.1, if the moments method is used consistently (in both data extraction and fitting), none of the above issues arise and the fitted parameters are recovered correctly.

6.3 Choosing an optimizer

In our tests, the NS optimizer appears to be more robust than the LMA and NG optimizers, but this does not necessarily make it the best choice.

The LMA fails to converge to the right solution for 2 per cent of the mock observations from cases where it gets trapped in a local minimum. Due to the fast convergence properties of the algorithm, we can achieve a $\sim$ 100 per cent success rate by repeating the fitting procedure (using better initial estimates for the model parameters) for this 2 per cent in only a few seconds. However, as mentioned in Section 4.1, the majority of the failed fits occur for highly

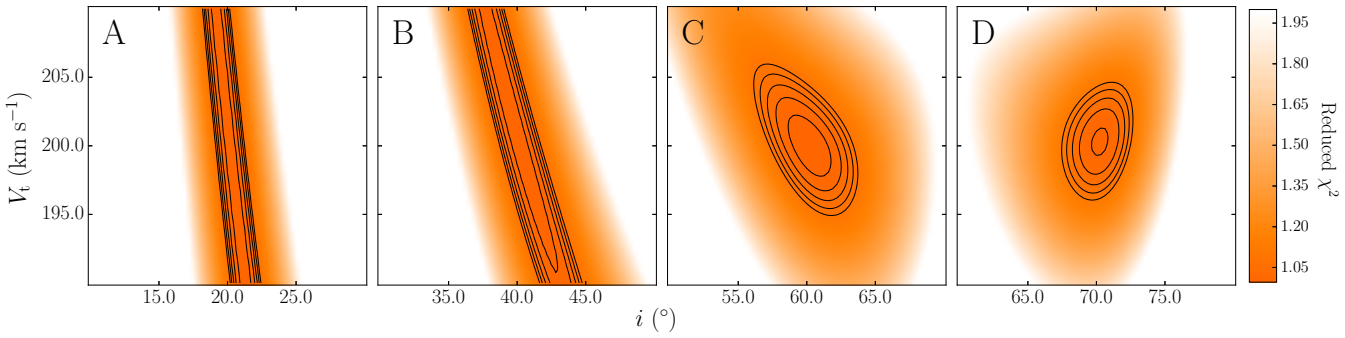

Figure 12. Reduced χ^2 space as a function of disc inclination (i) and asymptotic circular velocity (V_t). The contours correspond to 1, 2, 3, 4, and 5σ confidence intervals. Panels A, B, C, D correspond to galaxies with $V_t = 200$ km s^{-1} and $i = 20°$, 40°, 60°, and 70° respectively.

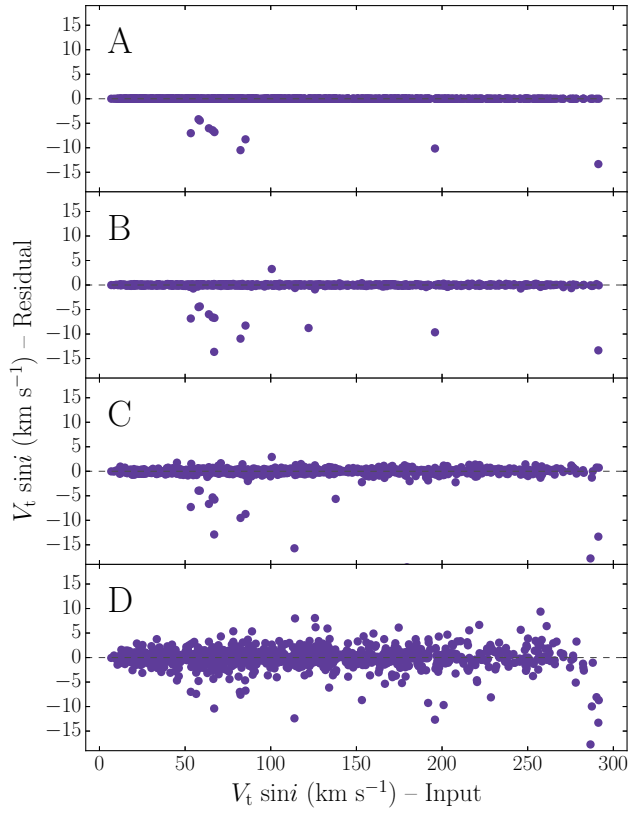

Figure 13. Fitting results of the mock data with different levels of noise using the Levenberg–Marquardt Algorithm. We plot the residuals of the fitted $V_t \sin i$ products as a function of the input parameters. The rows from top to bottom correspond to datasets with 0, 1, 5, and 20 km s^{-1} noise on the velocity map, and 0, 1.4, 7, and 28 km s^{-1} noise on the velocity dispersion map. The error bars are omitted for the sake of clarity.

inclined galaxies, for which a thin disc model is not a good representation anyway.

The LMA's fast convergence properties can also help to tackle one of its weaknesses: parameter error estimation. The derived error of the fitted parameters may not be robust since they are based on an elliptical Gaussian approximation for the likelihood at the minimum. To overcome this issue, one can perform further χ^2 analysis around the best fit, or perform a Monte Carlo simulation by repeating the fitting procedure multiple times using different noise realizations.

On the contrary, the NS optimizer does not suffer from any of the above issues because, instead of blindly following the χ^2 gradient or curvature of the parameter space, it explores the likelihood of the posterior probability distribution thoroughly but efficiently. This make NS less likely to converge to a wrong solution and also allows for robust error estimation. Indeed, in our tests with zero-noise mock data, the NS optimizer was never trapped into a local optimum and achieved a ~ 100 per cent success rate. However, the above advantages come with the price of higher computational cost. The NS algorithm requires two orders of magnitude more iterations compared to the LMA.

Lastly, the NG optimizer was also able to fit the majority of the galaxies in the zero-noise mock dataset with great success. However, a direct comparison with the other two optimizers would not be fair because only four of the model parameters were left free when fitting with the NG optimizer. Reducing the number of free parameters in our model was necessary in order to perform the fitting procedure in a reasonable amount of time. Despite its inefficiency, as already mentioned in Section 2.3.3, the NG optimizer is trivial to implement as a first solution when other optimization algorithms are not available. The time required to develop the NG code can be a few orders of magnitude less than the time required to develop sophisticated optimization software such as the MPFIT and MULTINEST libraries.

When it comes to choosing an optimization technique, there is no such thing as 'one-size-fits-all'. When fitting simple models that are unlikely to cause multimodalities on the parameter space, the LMA optimizer should suffice. In case we fit complex models with a large number of parameters that may result multimodalities or strong degeneracies, the NS optimizer will most likely give better results.

6.4 Accelerating science

By accelerating the computations involved in galaxy kinematic modelling using the GPU we can increase the pace of scientific discovery. Besides the obvious benefit of analysing results faster, the GPU may benefit kinematic modelling in several other important ways.

By providing much shorter computation times it is now feasible to fit multiple models to the data in the same com-

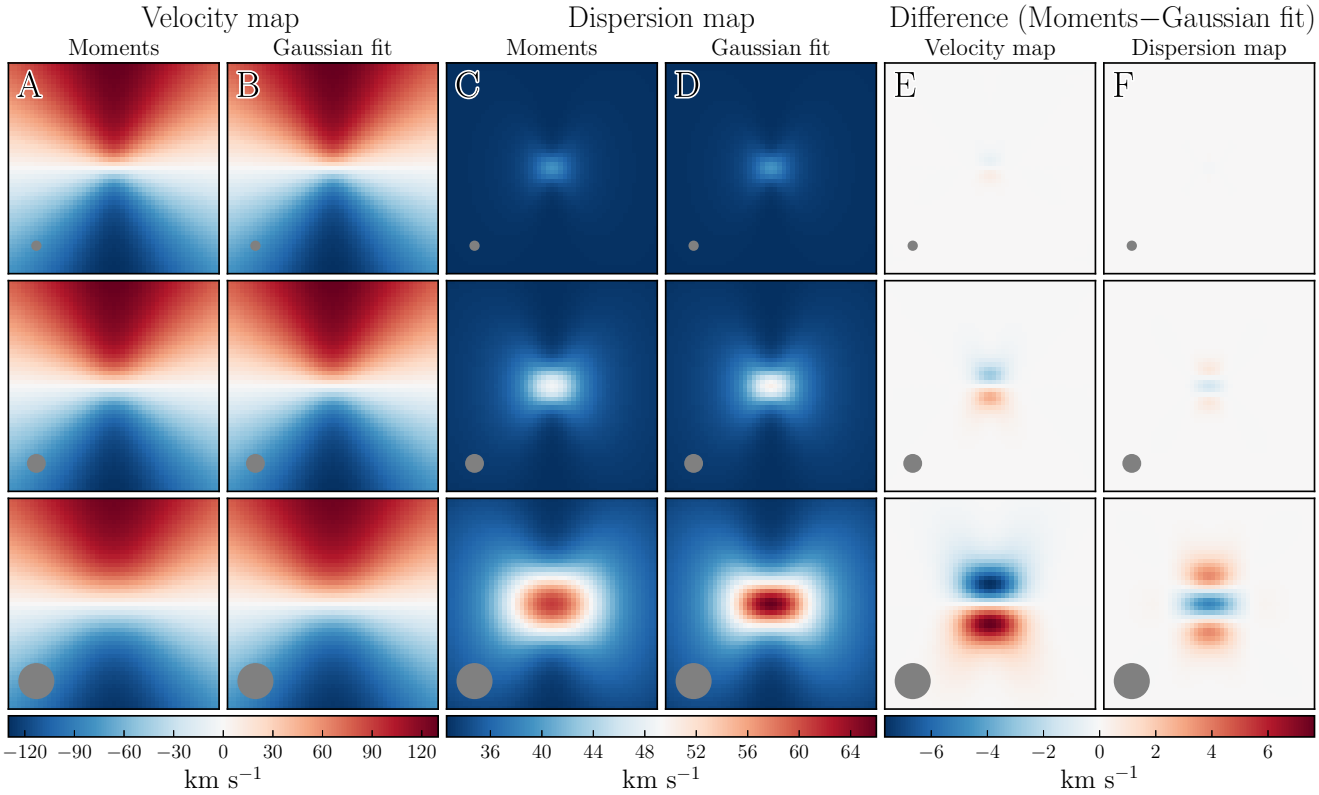

Figure 14. Comparison of two different methods for extracting velocity and velocity dispersion maps from intensity datacubes. (A, C) Velocity and velocity dispersion maps extracted using the first- and second-order moments of the cube's spectral lines. (B, D) Velocity and velocity dispersion maps extracted by fitting Gaussian profiles to the cube's spectral lines. (E, F) The differences in the velocity and velocity dispersion maps between the two methods. The seeing FWHM from top to bottom corresponds to 2, 4, and 8 pixels respectively. The spatial size of the data is 49×49 pixels and the velocity resolution is 30 km s^{-1}. The scale radius is 9 pixels, the turn-over radius is 3 pixels, the disc inclination is 45°, the maximum rotation velocity is 200 km s^{-1}, and the intrinsic velocity dispersion is 30 km s^{-1}.

pute time. This can be beneficial when no a priori information about the kinematics of the galaxy is available. Alternatively, one could fit a very complicated model using the GPU in the same amount of time that one could fit a simpler model using the CPU. Using more complicated models can eventually lead to a better understanding of the kinematics of the galaxy and potentially provide better constraints on galaxy's physical quantities.

It is known that the majority of the optimization methods used for model fitting in astronomy are not global optimizers but local ones. As a result there is a chance that the optimizer will not converge to the global optimal solution leading to the wrong results and eventually to the wrong science conclusions. This can happen for different reasons including parameter space multimodalities, low resolution data and bad initial model parameter values. With the acceleration the GPU provides, one can now fit the same model to the same data multiple times using either different optimizers or the same optimizer but starting from a different position in the parameter space. This allows us to cross-check the fitting results before proceeding to further analysis.

The GPU can also benefit the process of calculating uncertainties in the model parameters. This task usually requires the additional model evaluations, or even the repetition of the fitting procedure, for the same object hundreds or thousands of times.

Lastly, the use of the GPU could potentially enable different approaches to galaxy kinematic modelling, such as interactive visualization-directed model fitting which allows the user to build and fit complex kinematic models interactively while exploring the data (Fluke et al. 2010), in near real-time. In addition, it could potentially allow the user to provide better model parameter priors to the optimization algorithm. This approach to modelling can be useful when the observed galaxies feature exotic or unique kinematic structures that cannot be described by any of the models available in the automated modelling procedure. Although it may be demanding in terms of human interaction, it will allow scientist to model selected galaxies in greater detail. Steffen et al. (2011) developed a similar software which allows the user to model planetary nebulae interactively, and mentioned the need for GPU acceleration in such tools.

6.5 Fitting large-scale surveys

To obtain a first glimpse of the degree of speedup the graphics processing unit has to offer when fitting large-scale surveys, we estimate the computational time that is required to fit the entire SAMI (Bryant et al. 2015) and MaNGA (Bundy et al. 2015) surveys. We do so by fitting simulated

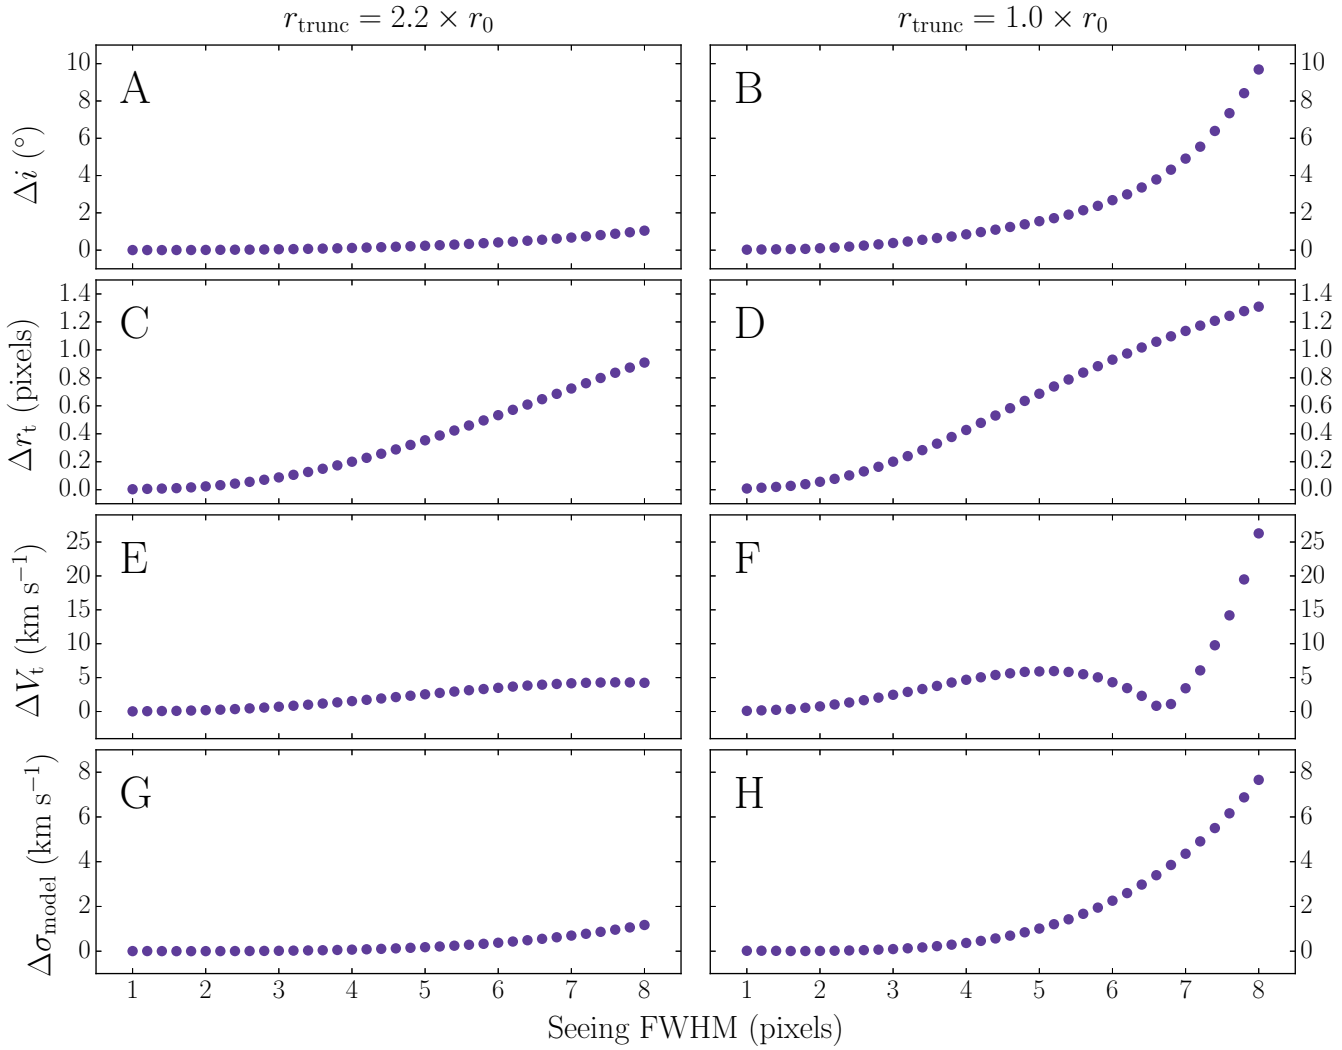

Figure 15. The error between the actual and fitted parameters as a function of seeing FWHM. We fit a model that uses the spectral moments approach to data that have been extracted using Gaussian profile fitting. The rows from top to bottom correspond to the error in the inclination, turn-over radius, circular velocity and intrinsic velocity dispersion respectively. The scale radius is 9 pixels, the turn-over radius is 3 pixels, the disc inclination is 45°, the maximum rotation velocity is 200 km s^{-1}, and the intrinsic velocity dispersion is 30 km s^{-1}. The left column corresponds to data truncated at 2.2 × exponential disc scale lengths, while the right column corresponds to data truncated at 1.0 × exponential disc scale lengths.

observations that match the properties of the real survey data.

The simulated SAMI observations have a spatial sampling of 0.5 arcsec pixel^{-1} across a field of view of ∼ 15 arcsec. The typical size of the point spread function is ∼ 2.1 arcsec, while the velocity resolution is ∼ 67 km s^{-1}. This configuration, along with an upsampling factor of $u = 2$ for all three dimensions, result in an intensity cube model of 128 × 128 × 256 pixels.

The simulated MaNGA observations have a spatial sampling of 0.5 arcsec pixel^{-1}, but the field of view can vary depending on the IFU size. MaNGA uses six different IFU configurations of 7, 19, 37, 61, 91 and 127 fibres, with the first one used for flux calibration. Each fibre has a diameter of 2 arcsec resulting in IFUs with diameters between 12 and 32 arcsec on-sky. We assume that the number of galaxies observed is uniformly distributed among the five different IFU configurations (2000 galaxies with each IFU setup). The typical size of the point spread function is ∼ 2.5 arcsec, while the velocity resolution is ∼ 120 km s^{-1}. This configuration, along with an upsampling factor of $u = 2$ for all three dimensions, result in intensity cube models with spatial dimensions of 128 × 128 and 256 × 256 pixels, and a spectral dimension of 128 pixels.

We investigate the execution time for all four hardware configurations using the Levenberg–Marquardt and Nested Sampling optimizers. For the LMA scenario we assume that a mean of ∼ 200 iterations is required before converging to a solution, while for the NS optimizer we assume a mean of ∼ 20000 iterations.

The execution time for the two surveys is shown in Fig. 16. It clearly shows the improvement achieved by the GPU implementation. While a single-threaded CPU implementation would only allow for analysis using the LMA op-

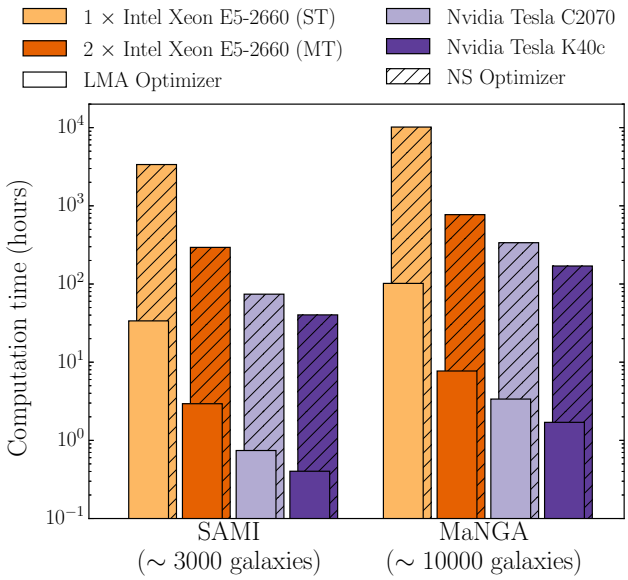

Figure 16. Computational time across different hardware configurations and optimization methods for fitting SAMI and MaNGA surveys. The computation time does not include the time spent on the optimizer's logic.

timizer in a reasonable amount of time ($\lesssim$ 200 hours), such performance can be reached for any optimizer using GPU programming. However, the speed achieved using the LMA optimizer and the GPU ($\lesssim$ 2 hours) shows that more sophisticated models could be developed. Moreover, if further decrease of the computation time is desired, astronomers may have to shift from the traditional approach of analysing data on the desktop, to the use of remote services and resources in the form of web applications and GPU-enabled supercomputers.

6.6 Next steps

The substantial speedup gained by using the graphics processing unit opens up new directions in galaxy kinematic modelling. It allows us to perform analysis that was hard to carry out before due to the enormous computational requirements. In future work we will investigate:

(1) The implementation of more complex methods and kinematic models that would previously take a prohibitive amount of time to fit on the CPU. These could include methods to fit directly to datacubes and models that account for non-circular or asymmetric motions, disc warps and multiple galactic components.
(2) Interactive and visualization-directed approaches to galaxy kinematic modelling. This will allow scientists to build complex galaxy models, featuring unusual and peculiar kinematic structures, while exploring their data.
(3) Distributed kinematic analysis of large-scale 3D spectroscopic surveys using remote HPC facilities and web services.

An implementation of a gradient- and curvature-based optimizer, such as the Levenberg–Marquardt Algorithm, running solely on the GPU would also be beneficial. This would allow the extraction of the velocity and velocity dispersion maps from intensity cubes with complex spectra, where the first- and second-order moments approach may fail (e.g., multiple emission and absorption lines, asymmetric continuum etc.).

Moreover, we believe that the majority of kinematic modelling software and tools, such as TIRIFIC, ^{3D}BAROLO, GALPAK^{3D}, DISKFIT, KINEMETRY and others, can benefit from a GPU implementation because they all perform a high number of independent pixel operations on two- and three-dimensional data.

As part of this work we have developed a high-performance software tool for galaxy kinematic modelling: GBKFIT. The software along with its source code, documentation and other supporting material is available at: http://supercomputing.swin.edu.au/gbkfit. For more information about the development and implementation of GBKFIT, see Appendix B.

7 SUMMARY AND CONCLUSION

We have investigated the use of the graphics processing unit (GPU) as an accelerator for the computationally demanding fitting procedure in the context of galaxy kinematic modelling.

For our analysis we used three different optimization techniques including the Levenberg–Marquardt algorithm, Nested Sampling, and a naive brute-force technique based on Nested Grids.

To assess the consistency of our method we fitted kinematic models to data from three different sources including artificially generated observations and observations from the low redshift surveys GHASP and DYNAMO. The last two were chosen as representatives of high- and low-spatial-resolution surveys respectively. For all three datasets only a minimal fine-tuning of our fitting code was required.

We found that our method can successfully recover the parameters of zero-noise mock galaxies with a success rate up to $\sim$ 100 per cent. In addition, our method was able to fit mock observations with low and high levels of noise with great success: $\sim$ 68 and $\sim$ 95 per cent of the the galaxy parameters are enclosed within the 1 and 2σ confidence intervals of our fits respectively. Furthermore, we were able to verify the kinematic modelling results of the GHASP and DYNAMO surveys by comparing our results to the literature.

For our performance tests we used four different hardware configurations: two high-end CPUs in single-threaded and multi-threaded mode, a low-end GPU and a high-end GPU configuration.

Our tests and analysis led to the following findings:

(1) The GPU can substantially accelerate most of the steps involved in the model fitting procedure. We found that our GPU code can perform up to $\sim$ 100 times faster than a single-threaded CPU implementation, and up to $\sim$ 10 times faster than a multi-threaded dual-CPU solution.
(2) Taking advantage of the massive processing power of the GPU allows us to use simpler, naive and less efficient approaches which are conceptually easier to implement.
(3) The substantial acceleration the GPU provides allows

us to perform additional analysis on our data. This can include the fitting of multiple models to the same data or the use of multiple optimization methods to verify results.

(4) Although to achieve the highest possible performance on a GPU might require a high amount of expertise, we can still achieve substantial performance improvement (compared to the CPU) by following a few basic GPU programming guidelines.

(5) Caution is advised when fitting high-resolution data since the memory available on GPU is much less than the main system memory.

We are poised on the brink of a revolution of our understanding of galaxies in the nearby universe. The currently underway or upcoming large-scale 3D spectroscopic surveys (e.g., SAMI, MaNGA, WALLABY) will reveal the internal structure and kinematics of several thousands of galaxies. This vast amount of data will push the computational requirements for analysis up by several orders of magnitude. We expect that massively parallel architectures, such as the graphics processing unit, will play a crucial role in analysing and understanding this fire hose of new data, which will eventually lead to new scientific discoveries.

ACKNOWLEDGEMENTS

We thank Andrew Green for providing us with the data of the DYNAMO survey. G.B. thanks Sylvia Mackie for providing linguistic advice. We are grateful for the constructive comments of the referee, which have helped to improve this paper. This work was performed on the gSTAR national facility at Swinburne University of Technology. gSTAR is funded by Swinburne and the Australian Government's Education Investment Fund. This research has made use of the Fabry Perot database, operated at CeSAM/LAM, Marseille, France. We acknowledge the usage of the HyperLeda database (`http://leda.univ-lyon1.fr`). This research was supported under Australian Research Council's Discovery Projects funding scheme (project number DP130101460).

REFERENCES

Adelman-McCarthy J. K., et al., 2006, ApJS, 162, 38
Anthony S. M., Granick S., 2009, Langmuir, 25, 8152
Ashman K. M., 1992, PASP, 104, 1109
Bacon R., et al., 1995, A&AS, 113, 347
Bacon R., et al., 2015, A&A, 575, A75
Barbosa F. K. B., Storchi-Bergmann T., Cid Fernandes R., Winge C., Schmitt H., 2006, MNRAS, 371, 170
Barsdell B. R., Barnes D. G., Fluke C. J., 2010, in Mizumoto Y., Morita K.-I., Ohishi M., eds, Astronomical Society of the Pacific Conference Series Vol. 434, Astronomical Data Analysis Software and Systems XIX. p. 209 (arXiv:1001.2048)
Begeman K. G., 1987, PhD thesis, Kapteyn Institute, (1987)
Belleman R. G., Bédorf J., Portegies Zwart S. F., 2008, New Astron., 13, 103
Bershady M. A., Verheijen M. A. W., Swaters R. A., Andersen D. R., Westfall K. B., Martinsson T., 2010, ApJ, 716, 198
Blake C., et al., 2012, MNRAS, 425, 405
Bland-Hawthorn J., 2015, in Ziegler B. L., Combes F., Dannerbauer H., Verdugo M., eds, IAU Symposium Vol. 309, IAU Symposium. pp 21–28 (arXiv:1410.3838), doi:10.1017/S1743921314009247
Bosma A., 1978, PhD thesis, Groningen Univ., (1978)
Bosma A., 1981, AJ, 86, 1791
Bouché N., Carfantan H., Schroetter I., Michel-Dansac L., Contini T., 2015, AJ, 150, 92
Brandon L., Boyd C., Govindaraju N., 2008, in Multimedia and Expo, 2008 IEEE International Conference on. pp 5–8, doi:10.1109/ICME.2008.4607357
Briggs F. H., 1990, ApJ, 352, 15
Brodtkorb A. R., Hagen T. R., Sætra M. L., 2013, JPDC, 73, 4
Broeils A. H., Courteau S., 1997, in Persic M., Salucci P., eds, Astronomical Society of the Pacific Conference Series Vol. 117, Dark and Visible Matter in Galaxies and Cosmological Implications. p. 74 (arXiv:astro-ph/9610264)
Bryant J. J., et al., 2015, MNRAS, 447, 2857
Budruk R., Anderson D., Shanley T., 2004, PCI express system architecture. Addison-Wesley Professional
Bundy K., et al., 2015, ApJ, 798, 7
Cameron E., Pettitt A., 2014, Statist. Sci., 29, 397
Cao J., Novstrup K. A., Goyal A., Midkiff S. P., Caruthers J. M., 2009, in Proceedings of the 23rd International Conference on Supercomputing. ICS '09. ACM, New York, NY, USA, pp 450–459, doi:10.1145/1542275.1542338, http://doi.acm.org/10.1145/1542275.1542338
Cappellari M., Emsellem E., 2004, PASP, 116, 138
Cappellari M., et al., 2011, MNRAS, 413, 813
Coccato L., Swaters R. A., Rubin V. C., D'Odorico S., McGaugh S. S., 2008, A&A, 490, 589
Contini T., et al., 2012, A&A, 539, A91
Cooray A., Sheth R., 2002, Phys. Rep., 372, 1
Courteau S., 1997, AJ, 114, 2402
Cresci G., et al., 2009, ApJ, 697, 115
Dagum L., Menon R., 1998, Computational Science Engineering, IEEE, 5, 46
Davies R. L., et al., 2001, ApJ, 548, L33
Davies R., et al., 2011, ApJ, 741, 69
Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371
Di Teodoro E. M., Fraternali F., 2015, MNRAS, 451, 3021
Duhamel P., Vetterli M., 1990, Signal Processing, 19, 259
Efstathiou G., Sutherland W. J., Maddox S. J., 1990, Nature, 348, 705
Eichenberger A. E., Wu P., O'Brien K., 2004, SIGPLAN Not., 39, 82
Epinat B., et al., 2008a, MNRAS, 388, 500
Epinat B., Amram P., Marcelin M., 2008b, MNRAS, 390, 466
Epinat B., Amram P., Balkowski C., Marcelin M., 2010, MNRAS, 401, 2113
Epinat B., et al., 2012, A&A, 539, A92
Erb D. K., Steidel C. C., Shapley A. E., Pettini M., Reddy N. A., Adelberger K. L., 2006, ApJ, 646, 107
Evans I. N., Koratkar A. P., Storchi-Bergmann T., Kirkpatrick H., Heckman T. M., Wilson A. S., 1996, ApJS, 105, 93
Falcón-Barroso J., et al., 2004, MNRAS, 350, 35
Fatica M., 2009, in Proceedings of 2Nd Workshop on General Purpose Processing on Graphics Processing Units. GPGPU-2. ACM, New York, NY, USA, pp 46–51, doi:10.1145/1513895.1513901, http://doi.acm.org.ezproxy.lib.swin.edu.au/10.1145/1513895.151390
Feroz F., Hobson M. P., 2008, MNRAS, 384, 449
Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601
Fluke C. J., English J., Barnes D. G., 2010, in Mizumoto Y., Morita K.-I., Ohishi M., eds, Astronomical Society of the Pacific Conference Series Vol. 434, Astronomical Data Analysis Software and Systems XIX. p. 167 (arXiv:1001.2043)
Fluke C. J., Barnes D. G., Barsdell B. R., Hassan A. H., 2011, Publ. Astron. Soc. Australia, 28, 15
Förster Schreiber N. M., et al., 2006, ApJ, 645, 1062
Förster Schreiber N. M., et al., 2009, ApJ, 706, 1364

Freeman K. C., 1970, ApJ, 160, 811
Frigo M., Johnson S., 2005, Proceedings of the IEEE, 93, 216
García-Ruiz I., Sancisi R., Kuijken K., 2002, A&A, 394, 769
Garrido O., Marcelin M., Amram P., 2004, MNRAS, 349, 225
Genzel R., et al., 2008, ApJ, 687, 59
Geyer C. J., 1991, in Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface. Interface Foundation of North America
Gnedin O. Y., Weinberg D. H., Pizagno J., Prada F., Rix H.-W., 2007, ApJ, 671, 1115
Gnerucci A., et al., 2011, A&A, 528, A88
Gratadour D., et al., 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. p. 6, doi:10.1117/12.2056358
Green A. W., et al., 2014, MNRAS, 437, 1070
Haario H., Saksman E., Tamminen J., 2001, Bernoulli, pp 223–242
Hassan A. H., Fluke C. J., Barnes D. G., Kilborn V. A., 2013, MNRAS, 429, 2442
Hastings W. K., 1970, Biometrika, 57, 97
Hillis W. D., Steele Jr G. L., 1986, Communications of the ACM, 29, 1170
Józsa G. I. G., Kenn F., Klein U., Oosterloo T. A., 2007, A&A, 468, 731
Kaczmarski K., Rzazewski P., 2013, in Pechenizkiy M., Wojciechowski M., eds, Advances in Intelligent Systems and Computing, Vol. 185, New Trends in Databases and Information Systems. Springer Berlin Heidelberg, pp 37–46, doi:10.1007/978-3-642-32518-2˙4, http://dx.doi.org/10.1007/978-3-642-32518-2_4
Kannappan S. J., 2001, PhD thesis, Harvard University
Kannappan S. J., Fabricant D. G., Franx M., 2002, AJ, 123, 2358
Kauker D., Sanftmann H., Frey S., Ertl T., 2010, in Computer and Information Technology (CIT), 2010 IEEE 10th International Conference on. pp 1152–1157, doi:10.1109/CIT.2010.209
Kenworthy M. A., Parry I. R., Taylor K., 2001, PASP, 113, 215
Koribalski B. S., 2012, Publ. Astron. Soc. Australia, 29, 359
Krajnović D., Cappellari M., de Zeeuw P. T., Copin Y., 2006, MNRAS, 366, 787
Kurihara K., Chaiken D., Agarwal A., 1991, in In Proceedings International Symposium on Shared Memory Multiprocessing. IPS Press, pp 91–101
Landaverde R., Zhang T., Coskun A., Herbordt M., 2014, in High Performance Extreme Computing Conference (HPEC), 2014 IEEE. pp 1–6, doi:10.1109/HPEC.2014.7040988
Law D. R., Steidel C. C., Erb D. K., Larkin J. E., Pettini M., Shapley A. E., Wright S. A., 2009, ApJ, 697, 2057
Levenberg K., 1944, Quarterly Journal of Applied Mathmatics, II, 164
Lewis A., Bridle S., 2002, Phys. Rev. D, 66, 103511
Markwardt C. B., 2009, in Bohlender D. A., Durand D., Dowler P., eds, Astronomical Society of the Pacific Conference Series Vol. 411, Astronomical Data Analysis Software and Systems XVIII. p. 251 (arXiv:0902.2850)
Marquardt D. W., 1963, SIAM Journal on Applied Mathematics, 11, 431
Mathewson D. S., Ford V. L., Buchhorn M., 1992, ApJS, 81, 413
Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, The journal of chemical physics, 21, 1087
Mo H., van den Bosch F. C., White S., 2010, Galaxy Formation and Evolution. Cambridge University Press
Mukherjee P., Parkinson D., Liddle A. R., 2006, ApJ, 638, L51
Paturel G., Petit C., Prugniel P., Theureau G., Rousseau J., Brouty M., Dubois P., Cambrésy L., 2003, A&A, 412, 45
Perret V., Renaud F., Epinat B., Amram P., Bournaud F., Contini T., Teyssier R., Lambert J.-C., 2014, A&A, 562, A1
Planck Collaboration et al., 2015, preprint, (arXiv:1502.01589)
Roberts M. S., Rots A. H., 1973, A&A, 26, 483
Rogstad D. H., Shostak G. S., 1972, ApJ, 176, 315
Rogstad D. H., Lockhart I. A., Wright M. C. H., 1974, ApJ, 193, 309
Rubin V. C., Thonnard N., Ford Jr. W. K., 1978, ApJ, 225, L107
Ryoo S., Rodrigues C. I., Baghsorkhi S. S., Stone S. S., Kirk D. B., Hwu W.-m. W., 2008, in Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming. PPoPP '08. ACM, New York, NY, USA, pp 73–82, doi:10.1145/1345206.1345220, http://doi.acm.org/10.1145/1345206.1345220
Sánchez S. F., et al., 2012, A&A, 538, A8
Sarzi M., et al., 2006, MNRAS, 366, 1151
Sharp R., et al., 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. (arXiv:astro-ph/0606137), doi:10.1117/12.671022
Shaw J. R., Bridges M., Hobson M. P., 2007, MNRAS, 378, 1365
Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V., eds, American Institute of Physics Conference Series Vol. 735, American Institute of Physics Conference Series. pp 395–405, doi:10.1063/1.1835238
Sofue Y., Rubin V., 2001, ARA&A, 39, 137
Spano M., Marcelin M., Amram P., Carignan C., Epinat B., Hernandez O., 2008, MNRAS, 383, 297
Spekkens K., Sellwood J. A., 2007, ApJ, 664, 204
Steffen W., Koning N., Wenger S., Morisset C., Magnor M., 2011, IEEE Transactions on Visualization and Computer Graphics, Volume 17, Issue 4, p.454-465, 17, 454
Stone R. C., 1989, AJ, 97, 1227
Thomas D. B., Howes L., Luk W., 2009, in Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays. FPGA '09. ACM, New York, NY, USA, pp 63–72, doi:10.1145/1508128.1508139, http://doi.acm.org/10.1145/1508128.1508139
Thompson A. C., Fluke C. J., Barnes D. G., Barsdell B. R., 2010, New Astron., 15, 16
Trimble V., 1987, ARA&A, 25, 425
Tsvetanov Z. I., Petrosian A. R., 1995, ApJS, 101, 287
Tully R. B., Fisher J. R., 1977, A&A, 54, 661
Verolme E. K., et al., 2002, MNRAS, 335, 517
Wayth R. B., Greenhill L. J., Briggs F. H., 2009, PASP, 121, 857
Weiner B. J., et al., 2006, ApJ, 653, 1049
White S. D. M., Rees M. J., 1978, MNRAS, 183, 341
Wiegert T., English J., 2014, New Astron., 26, 40
Wienke S., Springer P., Terboven C., an Mey D., 2012, in Kaklamanis C., Papatheodorou T., Spirakis P., eds, Lecture Notes in Computer Science, Vol. 7484, Euro-Par 2012 Parallel Processing. Springer Berlin Heidelberg, pp 859–870, doi:10.1007/978-3-642-32820-6˙85, http://dx.doi.org/10.1007/978-3-642-32820-6_85
Wisnioski E., et al., 2011, MNRAS, 417, 2601
Wisnioski E., et al., 2015, ApJ, 799, 209
Wright S. A., et al., 2007, ApJ, 658, 78
Wright S. A., Larkin J. E., Law D. R., Steidel C. C., Shapley A. E., Erb D. K., 2009, ApJ, 699, 421
Yang Y., et al., 2008, A&A, 477, 789
Zhu X., Zhang D., 2013, PLoS ONE, 8, 76665
Zwicky F., 1933, Helvetica Physica Acta, 6, 110
de Rijcke S., Michielsen D., Dejonghe H., Zeilinger W. W., Hau G. K. T., 2005, A&A, 438, 491
de Vaucouleurs G., 1959, Handbuch der Physik, 53, 311

APPENDIX A: USING AN MCMC OPTIMIZER

Bayesian model parameter estimation through Markov Chain Monte Carlo (MCMC) techniques has been used extensively in the field of astronomy (Lewis & Bridle 2002;

Wisnioski et al. 2011; Blake et al. 2012; Bouché et al. 2015). The main strength of the MCMC techniques is their ability to explore the parameter space efficiently, and provide robust estimates of the posterior probability distribution of the model parameters.

Standard MCMC methods, such as the Metropolis–Hastings algorithm (M–H, Hastings 1970), can be problematic or experience slow convergence when the likelihood landscape of the parameter space features multiple modes or strong degeneracies. In addition, simple and easy-to-implement MCMC techniques, such as the Random-Walk Metropolis algorithm (RWM, Metropolis et al. 1953), can be difficult to use because they require manual fine tuning of the proposal distribution in order to achieve high sampling efficiency.

Over the years several MCMC algorithms have been developed to address the above issues. Adaptive Metropolis (AM, Haario et al. 2001) is an extension to the RWM and adapts its proposals based on the history of the chain, thus no manual tuning of the proposal distribution is required. Metropolis-Coupled MCMC (MC^3, Geyer 1991) is better able to approximate multimodal distributions by running multiple chains of different temperatures.

Unfortunately only a few MCMC codes are publicly available and in a form of reusable libraries compatible with our software (i.e., C/C++ libraries). The majority of them are simple implementations of the standard M–H and RWM algorithms, and suffer from the issues mentioned above. In our case, while multimodalities are rarely formed on the likelihood landscape of the parameter space, the inclination and circular velocity parameters of our models will always form a strong degeneracy for galaxies close to face-on.

We used the publicly available GBPMCMC[10] library which implements a RWM sampler and provides an auto-tune facility to ease the tuning of the proposal distribution prior to the MCMC burn-in and integration phases. This library has been used successfully in Wisnioski et al. (2011) to model the kinematics of high redshift galaxies. In that work, to remove the degeneracy between the inclination (i) and the circular velocity (V_t), they simplified the kinematic model by using the product of $V_t \times \sin i$ as a single parameter.

After fitting several mock observations with both the MCMC and NS optimizers we concluded that the latter is superior for the following reasons:

(1) Although the GBPMCMC library provides a facility to ease the tuning of the proposal distribution, careful adjustment of this facility is still required in order to build a proposal distribution that will result in good sampling efficiency.

(2) With the MULTINEST library we were able to achieve high quality fits without any fine tuning of the algorithm and only by using the default configuration of the library.

(3) The auto-tune facility of GBPMCMC requires a few hundreds or thousands likelihood evaluations which can add to the total runtime of the fitting procedure.

(4) An initial set of model parameter values is required when using the MCMC optimizer, while this is not the case with the NS sampler.

(5) The MCMC optimizer requires a higher number of iterations in order to achieve a precision similar to NS.

(6) The MCMC optimizer does not check for the convergence of the chain during the runtime. The user has to perform convergence tests on the resulting chain after the integration phase is complete. On the other hand, NS convergence is checked during the runtime and it is controlled through a single criterion: *tol*.

(7) Finally, the MULTINEST library is modified compared to the original Nested Sampling algorithm specifically for sampling multimodal distributions and posteriors with pronounced degeneracies efficiently.

Since a 'vanilla' implementation of the RWM sampler has no advantage over the Nested Sampling implementation of the MULTINEST library, we decided not to use the MCMC optimizer in this work.

APPENDIX B: GBKFIT DEVELOPMENT AND IMPLEMENTATION NOTES

GBKFIT is an open source high-performance software for galaxy kinematic modelling which was developed as part of this work. It is written in C++ and supports all major operating systems such as Linux, Windows and Mac OS. To achieve this cross-platform compatibility we used free and open source tools, libraries and technologies while also conforming with the C++ coding standards and avoiding compiler-specific language extensions. To minimize the dependencies to external libraries we take advantage of C++ 11/14 features when possible. Furthermore, in order to provide a common unified compilation and installation procedure we make use of the CMake[11] build system.

GBKFIT is not a stand-alone executable but a library which can be linked with an executable and provide the necessary facilities for galaxy kinematic modelling. For this work, in order to communicate with the GBKFIT library and fit models to data we have developed a simple command-line application in C++.

The GBKFIT library features an extensible object-oriented architecture which allows the user to add support for arbitrary models and optimization techniques in the form of modules without modifying its source code. In addition, the modular architecture of GBKFIT allows the user to load new model and optimization modules during the runtime in the form of plugins without the need for recompilation.

So far we have developed the two kinematic models described in Section 2.1. Both models use CUDA in order to run on the GPU, but they also provide a multi-threaded CPU code path in case there is no Nvidia GPU available on the system. Furthermore, adapting GBKFIT to different Nvidia GPU models is done automatically by CUDA, which compiles the GPU code to binary code specifically for the underlying GPU architecture.

Finally, we have developed three optimizer modules which implement the three optimization techniques discussed in Section 2.3.

[10] http://github.com/gbpoole/gbpCode

[11] http://www.cmake.org/

APPENDIX C: FITTING RESULTS

Table C1. Kinematic properties of the GHASP sample (see Section 4.2).

Galaxy (UGC ID) (1)	$V_{\rm sys}$ (km s^{-1}) (2)	$V_{\rm sys,Epinat}$ (km s^{-1}) (3)	PA (°) (4)	$PA_{\rm Epinat}$ (°) (5)	i (°) (6)	$i_{\rm Epinat}$ (°) (7)	$V_{\rm max}$ (km s^{-1}) (8)	$V_{\rm max,Epinat}$ (km s^{-1}) (9)
IC476	4765 ± 1	4767 ± 3	73 ± 6	68 ± 6	55 ± 15	55 ± 24	107 ± 113	70 ± 22
IC2542	6111 ± 1	6111 ± 2	174 ± 2	174 ± 3	24 ± 11	20 ± 15	253 ± 127	290 ± 192
NGC5296	2255 ± 1	2254 ± 2	2 ± 3	2 ± 3	72 ± 10	65 ± 4	115 ± 101	80 ± 9
89	4513 ± 1	4510 ± 5	174 ± 2	177 ± 4	29 ± 8	33 ± 13	405 ± 155	343 ± 117
94	4546 ± 1	4548 ± 2	90 ± 1	94 ± 2	42 ± 3	42 ± 5	217 ± 12	209 ± 21
508	4637 ± 1	4641 ± 2	118 ± 2	123 ± 2	$25 \pm 7^*$	$25 \pm 7^*$	537 ± 27	553 ± 127
528	627 ± 1	628 ± 1	49 ± 1	52 ± 3	24 ± 6	21 ± 14	78 ± 32	84 ± 52
763	1147 ± 1	1148 ± 2	117 ± 1	117 ± 3	56 ± 3	54 ± 6	103 ± 3	104 ± 11
1117	-190 ± 1	-191 ± 1	198 ± 1	198 ± 3	$56 \pm 16^*$	$56 \pm 16^*$	92 ± 2	79 ± 17
1256	429 ± 1	428 ± 1	72 ± 1	73 ± 2	78 ± 1	76 ± 2	120 ± 9	105 ± 9
12632	412 ± 2	415 ± 2	44 ± 3	47 ± 4	$46 \pm 16^*$	$46 \pm 16^*$	143 ± 10	69 ± 19
12754	743 ± 1	742 ± 2	342 ± 1	342 ± 2	43 ± 4	53 ± 5	130 ± 8	123 ± 11
12893	1097 ± 1	1097 ± 2	260 ± 6	257 ± 5	16 ± 18	19 ± 19	84 ± 69	72 ± 67
1317	3092 ± 1	3090 ± 2	103 ± 1	106 ± 1	72 ± 1	73 ± 1	211 ± 4	205 ± 9
1437	4863 ± 1	4858 ± 2	301 ± 2	307 ± 2	43 ± 4	47 ± 4	229 ± 15	218 ± 15
1655	5432 ± 3	5427 ± 7	318 ± 5	318 ± 6	$45 \pm 18^*$	$45 \pm 18^*$	209 ± 16	205 ± 64
1736	1527 ± 1	1522 ± 2	27 ± 1	27 ± 3	49 ± 5	35 ± 14	206 ± 22	193 ± 68
1886	4841 ± 1	4836 ± 3	33 ± 1	35 ± 2	63 ± 2	62 ± 2	264 ± 6	267 ± 8
1913	537 ± 1	536 ± 2	288 ± 1	288 ± 3	47 ± 2	48 ± 9	111 ± 3	105 ± 16
2023	594 ± 1	593 ± 2	$315 \pm 9^*$	$315 \pm 9^*$	$19 \pm 18^*$	$19 \pm 18^*$	58 ± 9	53 ± 49
2034	567 ± 1	567 ± 2	$342 \pm 17^*$	$342 \pm 17^*$	$19 \pm 23^*$	$19 \pm 23^*$	38 ± 12	39 ± 46
2045	1523 ± 1	1525 ± 3	320 ± 1	319 ± 4	61 ± 1	61 ± 8	196 ± 4	185 ± 16
2080	893 ± 1	893 ± 1	336 ± 1	336 ± 2	$25 \pm 9^*$	$25 \pm 9^*$	123 ± 6	131 ± 42
2141	964 ± 1	965 ± 3	190 ± 1	191 ± 5	77 ± 4	74 ± 23	264 ± 80	157 ± 20
2183	1475 ± 1	1475 ± 4	158 ± 1	159 ± 3	49 ± 4	41 ± 10	136 ± 8	160 ± 32
2193	519 ± 1	519 ± 1	302 ± 2	305 ± 6	23 ± 10	6 ± 15	34 ± 7	174 ± 423
2455	368 ± 1	367 ± 2	$263 \pm 21^*$	$263 \pm 21^*$	$51 \pm 30^*$	$51 \pm 30^*$	31 ± 6	21 ± 12
2503	2372 ± 1	2377 ± 2	33 ± 1	34 ± 1	48 ± 4	49 ± 2	320 ± 12	285 ± 12
2800	1179 ± 2	1174 ± 2	294 ± 4	291 ± 4	38 ± 8	52 ± 13	122 ± 14	103 ± 20
2855	1190 ± 1	1188 ± 2	99 ± 1	100 ± 2	67 ± 1	68 ± 2	222 ± 7	229 ± 9
3013	2452 ± 1	2453 ± 6	193 ± 1	195 ± 4	$58 \pm 8^*$	$58 \pm 8^*$	215 ± 7	212 ± 21
3273	615 ± 1	615 ± 2	41 ± 2	42 ± 4	79 ± 3	82 ± 7	106 ± 18	106 ± 7
3334	3945 ± 1	3952 ± 13	275 ± 1	277 ± 6	$47 \pm 14^*$	$47 \pm 14^*$	378 ± 4	377 ± 85
3382	4488 ± 1	4490 ± 2	183 ± 2	184 ± 2	22 ± 11	18 ± 6	276 ± 180	335 ± 111
3429	899 ± 1	868 ± 4	327 ± 1	317 ± 3	55 ± 2	54 ± 8	243 ± 4	322 ± 30
3463	2680 ± 1	2679 ± 3	109 ± 1	110 ± 2	62 ± 1	63 ± 3	168 ± 2	168 ± 9
3521	4412 ± 1	4415 ± 2	255 ± 4	258 ± 3	62 ± 9	58 ± 5	171 ± 26	166 ± 12
3528	4337 ± 3	4340 ± 5	220 ± 3	223 ± 4	54 ± 8	42 ± 12	240 ± 31	276 ± 66
3574	1437 ± 1	1433 ± 1	101 ± 2	99 ± 3	54 ± 6	19 ± 10	74 ± 7	202 ± 96
3685	1796 ± 1	1795 ± 1	297 ± 3	298 ± 4	20 ± 8	12 ± 17	74 ± 15	133 ± 177
3691	2204 ± 1	2203 ± 2	249 ± 1	248 ± 3	63 ± 2	64 ± 4	151 ± 4	143 ± 10
3708	5165 ± 2	5161 ± 4	231 ± 2	230 ± 4	47 ± 11	44 ± 16	231 ± 198	234 ± 69
3709	5291 ± 2	5292 ± 4	232 ± 2	232 ± 2	62 ± 3	55 ± 4	251 ± 12	241 ± 14
3734	965 ± 1	966 ± 1	140 ± 1	139 ± 2	36 ± 6	43 ± 7	103 ± 21	108 ± 16
3740	2417 ± 1	2416 ± 2	247 ± 1	247 ± 4	55 ± 2	48 ± 14	84 ± 6	87 ± 20
3809	2201 ± 1	2200 ± 1	356 ± 1	357 ± 1	59 ± 2	58 ± 2	251 ± 4	258 ± 9
3826	1724 ± 1	1724 ± 2	255 ± 3	254 ± 5	5 ± 13	20 ± 19	266 ± 91	74 ± 66
3876	855 ± 1	854 ± 2	360 ± 2	358 ± 3	59 ± 3	59 ± 5	103 ± 4	112 ± 10
3915	4660 ± 1	4659 ± 3	30 ± 1	30 ± 2	45 ± 3	47 ± 4	205 ± 7	205 ± 16
4026	4890 ± 1	4892 ± 3	138 ± 1	139 ± 2	54 ± 2	56 ± 4	290 ± 6	284 ± 14
4165	505 ± 1	504 ± 1	266 ± 1	265 ± 2	42 ± 4	41 ± 10	71 ± 6	80 ± 18
4256	5256 ± 1	5252 ± 3	291 ± 3	291 ± 6	40 ± 15	38 ± 21	120 ± 48	123 ± 59
4273	2397 ± 1	2398 ± 2	211 ± 1	212 ± 2	57 ± 2	60 ± 4	194 ± 5	219 ± 11
4274	428 ± 1	430 ± 1	173 ± 2	175 ± 3	39 ± 5	27 ± 17	73 ± 7	102 ± 57
4284	535 ± 1	536 ± 2	176 ± 1	176 ± 3	69 ± 2	59 ± 9	167 ± 10	118 ± 14
4305	140 ± 1	139 ± 1	188 ± 1	189 ± 4	$40 \pm 27^*$	$40 \pm 27^*$	44 ± 3	48 ± 28
4325	508 ± 1	508 ± 2	57 ± 2	57 ± 3	52 ± 9	63 ± 14	97 ± 15	85 ± 13
4393	2122 ± 1	2119 ± 4	236 ± 5	250 ± 7	$50 \pm 9^*$	$50 \pm 9^*$	94 ± 9	47 ± 10
4422	4325 ± 1	4321 ± 2	36 ± 2	36 ± 2	21 ± 7	25 ± 8	410 ± 116	353 ± 94

Table C1 – *continued*

Galaxy (UGC ID) (1)	$V_{\rm sys}$ (km s^{-1}) (2)	$V_{\rm sys,Epinat}$ (km s^{-1}) (3)	PA (°) (4)	$PA_{\rm Epinat}$ (°) (5)	i (°) (6)	$i_{\rm Epinat}$ (°) (7)	$V_{\rm max}$ (km s^{-1}) (8)	$V_{\rm max,Epinat}$ (km s^{-1}) (9)
4456	5471 ± 1	5470 ± 1	125 ± 3	124 ± 3	43 ± 14	9 ± 14	51 ± 40	211 ± 321
4499	680 ± 1	682 ± 1	144 ± 4	141 ± 4	$50 \pm 14^*$	$50 \pm 14^*$	70 ± 8	62 ± 13
4543	1949 ± 1	1948 ± 2	319 ± 2	316 ± 4	66 ± 10	52 ± 15	79 ± 14	70 ± 15
4555	4233 ± 1	4235 ± 2	91 ± 2	90 ± 2	36 ± 9	38 ± 7	200 ± 137	185 ± 30
4770	7042 ± 3	7026 ± 3	280 ± 4	278 ± 3	32 ± 19	20 ± 13	290 ± 165	330 ± 194
4820	1349 ± 1	1350 ± 2	157 ± 1	157 ± 2	37 ± 2	38 ± 3	330 ± 10	336 ± 20
4936	1729 ± 1	1728 ± 1	295 ± 1	294 ± 2	11 ± 6	13 ± 12	292 ± 80	264 ± 227
5045	7664 ± 1	7667 ± 2	147 ± 1	148 ± 2	17 ± 6	16 ± 9	410 ± 118	429 ± 228
5175	3046 ± 1	3049 ± 2	145 ± 1	143 ± 2	56 ± 2	56 ± 3	189 ± 4	188 ± 10
5228	1869 ± 1	1869 ± 2	119 ± 1	120 ± 2	71 ± 1	72 ± 2	125 ± 5	125 ± 9
5251	1466 ± 1	1465 ± 3	260 ± 1	260 ± 3	74 ± 1	73 ± 6	126 ± 3	125 ± 9
5253	1323 ± 1	1322 ± 2	354 ± 1	356 ± 2	41 ± 2	40 ± 4	230 ± 8	235 ± 17
5316	1033 ± 1	1031 ± 2	131 ± 3	130 ± 3	76 ± 8	77 ± 4	205 ± 69	145 ± 9
5319	2441 ± 1	2439 ± 1	346 ± 2	345 ± 2	26 ± 8	30 ± 9	202 ± 68	180 ± 47
5373	291 ± 1	291 ± 2	58 ± 3	51 ± 8	85 ± 3	10 ± 18	27 ± 18	90 ± 162
5414	592 ± 1	592 ± 2	216 ± 3	219 ± 4	60 ± 13	71 ± 13	81 ± 22	74 ± 10
5510	1299 ± 1	1298 ± 2	200 ± 1	200 ± 3	22 ± 5	31 ± 10	216 ± 52	167 ± 44
5532	2802 ± 1	2802 ± 1	148 ± 1	147 ± 1	32 ± 2	32 ± 3	379 ± 15	398 ± 24
5721	528 ± 1	527 ± 6	$275 \pm 21^*$	$275 \pm 21^*$	$62 \pm 30^*$	$62 \pm 30^*$	102 ± 4	99 ± 29
5786	989 ± 1	992 ± 4	151 ± 1	153 ± 5	39 ± 3	53 ± 11	106 ± 6	80 ± 15
5789	729 ± 1	730 ± 2	28 ± 1	27 ± 3	70 ± 3	68 ± 10	124 ± 27	131 ± 10
5829	626 ± 1	626 ± 2	197 ± 3	198 ± 6	53 ± 10	34 ± 20	34 ± 16	48 ± 26
5840	582 ± 1	580 ± 1	333 ± 1	333 ± 3	18 ± 4	18 ± 11	253 ± 64	251 ± 138
5842	1245 ± 1	1245 ± 1	293 ± 2	292 ± 2	32 ± 9	47 ± 9	127 ± 62	115 ± 18
5931	1606 ± 1	1604 ± 3	358 ± 2	360 ± 5	56 ± 6	54 ± 16	171 ± 13	157 ± 32
5982	1572 ± 1	1573 ± 2	28 ± 1	28 ± 2	55 ± 2	55 ± 4	206 ± 6	199 ± 13
6118	1521 ± 1	1525 ± 2	339 ± 1	343 ± 3	$39 \pm 8^*$	$39 \pm 8^*$	135 ± 4	137 ± 24
6277	1194 ± 2	1191 ± 3	71 ± 3	76 ± 3	$17 \pm 17^*$	$17 \pm 17^*$	306 ± 47	270 ± 258
6419	1382 ± 1	1381 ± 2	38 ± 2	34 ± 7	22 ± 20	66 ± 19	76 ± 49	53 ± 11
6521	5841 ± 1	5842 ± 2	19 ± 1	20 ± 2	51 ± 3	46 ± 4	232 ± 7	249 ± 18
6523	5949 ± 1	5947 ± 2	352 ± 3	353 ± 3	$24 \pm 14^*$	$24 \pm 14^*$	117 ± 10	118 ± 63
6537	854 ± 1	856 ± 2	200 ± 1	200 ± 2	47 ± 1	47 ± 5	185 ± 2	187 ± 17
6628	863 ± 1	863 ± 1	183 ± 2	179 ± 2	$20 \pm 20^*$	$20 \pm 20^*$	201 ± 13	183 ± 168
6702	7334 ± 1	7332 ± 2	250 ± 2	256 ± 2	41 ± 6	38 ± 6	178 ± 31	195 ± 23
6778	951 ± 1	951 ± 1	343 ± 1	343 ± 2	50 ± 1	49 ± 4	232 ± 3	223 ± 14
6787	1153 ± 5	1157 ± 3	110 ± 1	112 ± 2	71 ± 2	70 ± 2	228 ± 11	232 ± 10
7021	1976 ± 1	1976 ± 3	264 ± 1	266 ± 2	$56 \pm 7^*$	$56 \pm 7^*$	190 ± 8	223 ± 18
7045	757 ± 1	758 ± 1	100 ± 1	99 ± 2	68 ± 1	68 ± 2	157 ± 2	160 ± 9
7154	1011 ± 1	1009 ± 1	276 ± 1	275 ± 2	67 ± 1	65 ± 3	157 ± 3	145 ± 9
7323	506 ± 1	505 ± 1	38 ± 1	38 ± 3	$51 \pm 11^*$	$51 \pm 11^*$	83 ± 2	84 ± 15
7766	807 ± 1	807 ± 1	322 ± 1	323 ± 2	68 ± 1	69 ± 3	118 ± 4	120 ± 9
7831	134 ± 1	136 ± 3	295 ± 1	290 ± 5	35 ± 3	56 ± 12	128 ± 7	92 ± 15
7853	529 ± 1	530 ± 2	216 ± 1	217 ± 4	$58 \pm 28^*$	$58 \pm 28^*$	116 ± 2	110 ± 35
7861	599 ± 1	598 ± 1	295 ± 3	297 ± 4	$47 \pm 24^*$	$47 \pm 24^*$	59 ± 7	50 ± 21
7876	944 ± 1	944 ± 1	343 ± 1	344 ± 3	45 ± 7	53 ± 9	118 ± 15	98 ± 14
7901	788 ± 1	788 ± 2	255 ± 1	254 ± 2	54 ± 1	53 ± 2	215 ± 3	215 ± 10
7971	458 ± 1	457 ± 2	30 ± 8	32 ± 11	50 ± 23	31 ± 27	26 ± 52	33 ± 27
7985	642 ± 1	642 ± 2	274 ± 1	276 ± 3	46 ± 2	49 ± 6	109 ± 2	112 ± 13
8334	484 ± 1	484 ± 1	99 ± 1	100 ± 1	67 ± 1	66 ± 1	229 ± 2	214 ± 9
8403	974 ± 1	975 ± 2	121 ± 1	121 ± 2	53 ± 2	57 ± 4	130 ± 2	128 ± 10
8490	191 ± 1	190 ± 1	168 ± 1	167 ± 3	43 ± 7	40 ± 15	86 ± 14	90 ± 29
8709	2406 ± 1	2405 ± 3	330 ± 1	330 ± 2	75 ± 1	76 ± 1	206 ± 2	207 ± 9
8852	2075 ± 1	2075 ± 1	62 ± 1	63 ± 2	50 ± 2	52 ± 3	192 ± 7	186 ± 10
8863	1782 ± 4	1789 ± 4	$218 \pm 7^*$	$218 \pm 7^*$	$77 \pm 13^*$	$77 \pm 13^*$	221 ± 21	191 ± 13
8898	3450 ± 1	3448 ± 2	31 ± 5	31 ± 6	36 ± 10	27 ± 20	47 ± 7	65 ± 45
8900	3511 ± 1	3511 ± 3	163 ± 1	161 ± 2	47 ± 5	57 ± 10	364 ± 28	345 ± 37
8937	2962 ± 1	2961 ± 5	185 ± 1	185 ± 3	33 ± 4	32 ± 12	317 ± 26	320 ± 105
9013	263 ± 1	262 ± 1	162 ± 2	164 ± 4	$21 \pm 16^*$	$21 \pm 16^*$	62 ± 3	62 ± 45
9179	295 ± 1	293 ± 2	52 ± 1	49 ± 4	30 ± 6	36 ± 14	129 ± 26	111 ± 36
9248	3878 ± 1	3865 ± 2	266 ± 2	261 ± 2	47 ± 6	58 ± 4	151 ± 16	166 ± 11
9358	1905 ± 1	1912 ± 3	186 ± 1	182 ± 2	57 ± 1	54 ± 4	213 ± 4	221 ± 14

Table C1 – *continued*

Galaxy (UGC ID) (1)	$V_{\rm sys}$ (km s^{-1}) (2)	$V_{\rm sys,Epinat}$ (km s^{-1}) (3)	PA (°) (4)	$PA_{\rm Epinat}$ (°) (5)	i (°) (6)	$i_{\rm Epinat}$ (°) (7)	$V_{\rm max}$ (km s^{-1}) (8)	$V_{\rm max,Epinat}$ (km s^{-1}) (9)
9363	1577 ± 1	1577 ± 1	147 ± 1	147 ± 3	$18 \pm 14^*$	$18 \pm 14^*$	142 ± 6	143 ± 105
9366	2109 ± 1	2109 ± 2	223 ± 1	225 ± 2	64 ± 1	62 ± 2	242 ± 2	241 ± 9
9406	2281 ± 1	2281 ± 2	131 ± 10	132 ± 12	51 ± 24	59 ± 25	25 ± 30	19 ± 10
9465	1486 ± 1	1485 ± 2	129 ± 1	127 ± 3	65 ± 2	65 ± 4	88 ± 3	97 ± 9
9576	1556 ± 1	1555 ± 2	120 ± 1	122 ± 3	38 ± 5	41 ± 11	120 ± 14	104 ± 25
9649	441 ± 1	440 ± 1	233 ± 1	235 ± 3	$54 \pm 6^*$	$54 \pm 6^*$	97 ± 3	94 ± 11
9736	3136 ± 1	3135 ± 2	217 ± 2	219 ± 2	48 ± 5	51 ± 5	201 ± 15	192 ± 16
9753	770 ± 1	764 ± 2	1 ± 1	3 ± 2	69 ± 1	69 ± 1	140 ± 2	138 ± 9
9858	2628 ± 2	2638 ± 3	67 ± 1	70 ± 2	75 ± 1	75 ± 2	169 ± 5	160 ± 9
9866	427 ± 1	430 ± 1	149 ± 1	148 ± 2	58 ± 3	56 ± 6	112 ± 9	114 ± 11
9943	1947 ± 1	1946 ± 1	266 ± 1	266 ± 2	58 ± 1	54 ± 2	201 ± 4	185 ± 10
9969	2521 ± 1	2516 ± 2	16 ± 1	16 ± 1	61 ± 2	61 ± 1	325 ± 6	311 ± 9
10075	830 ± 1	827 ± 1	210 ± 1	210 ± 1	61 ± 1	62 ± 2	176 ± 2	168 ± 9
10310	701 ± 2	702 ± 2	190 ± 4	187 ± 7	$42 \pm 20^*$	$42 \pm 20^*$	64 ± 7	66 ± 27
10359	910 ± 1	911 ± 2	281 ± 2	284 ± 3	$44 \pm 12^*$	$44 \pm 12^*$	127 ± 5	143 ± 30
10445	961 ± 1	961 ± 2	106 ± 2	110 ± 4	46 ± 3	47 ± 12	76 ± 4	77 ± 17
10470	1354 ± 1	1354 ± 2	287 ± 1	287 ± 2	$34 \pm 9^*$	$34 \pm 9^*$	161 ± 3	164 ± 39
10502	4290 ± 1	4291 ± 2	97 ± 2	99 ± 2	51 ± 7	50 ± 5	165 ± 15	163 ± 14
10521	834 ± 1	832 ± 2	20 ± 1	20 ± 2	59 ± 2	59 ± 3	121 ± 4	124 ± 9
10546	1265 ± 1	1268 ± 2	177 ± 2	182 ± 3	51 ± 6	42 ± 10	100 ± 8	106 ± 22
10564	1119 ± 1	1120 ± 2	149 ± 2	149 ± 3	75 ± 3	77 ± 6	102 ± 10	75 ± 8
10652	1090 ± 1	1089 ± 1	223 ± 1	225 ± 3	31 ± 8	21 ± 13	93 ± 39	141 ± 82
10757	1212 ± 1	1210 ± 2	59 ± 3	56 ± 6	39 ± 13	44 ± 22	88 ± 25	81 ± 33
10791	1313 ± 2	1318 ± 3	91 ± 6	92 ± 4	$34 \pm 20^*$	$34 \pm 20^*$	91 ± 18	95 ± 48
10897	1311 ± 1	1313 ± 1	113 ± 2	115 ± 3	$31 \pm 17^*$	$31 \pm 17^*$	111 ± 7	113 ± 56
11012	24 ± 1	25 ± 1	295 ± 1	299 ± 2	71 ± 1	72 ± 2	117 ± 2	117 ± 9
11124	1606 ± 1	1606 ± 2	180 ± 2	182 ± 3	49 ± 8	51 ± 10	98 ± 24	96 ± 15
11218	1476 ± 1	1477 ± 2	41 ± 1	42 ± 2	59 ± 1	58 ± 2	186 ± 2	185 ± 9
11269	2564 ± 4	2563 ± 6	276 ± 3	272 ± 3	66 ± 5	69 ± 4	270 ± 29	202 ± 13
11283	1945 ± 1	1944 ± 4	121 ± 2	120 ± 5	$34 \pm 17^*$	$34 \pm 17^*$	156 ± 5	173 ± 73
11300	481 ± 1	482 ± 1	168 ± 1	168 ± 2	66 ± 3	70 ± 3	117 ± 5	112 ± 9
11407	2405 ± 1	2402 ± 8	$65 \pm 10^*$	$65 \pm 10^*$	$64 \pm 22^*$	$64 \pm 22^*$	169 ± 5	158 ± 30
11429	4660 ± 2	4679 ± 7	208 ± 2	208 ± 6	$61 \pm 16^*$	$61 \pm 16^*$	224 ± 28	232 ± 35
11466	824 ± 1	826 ± 3	224 ± 1	226 ± 3	62 ± 2	66 ± 5	147 ± 4	133 ± 10
11470	3552 ± 4	3546 ± 5	45 ± 2	47 ± 3	47 ± 7	47 ± 7	386 ± 47	380 ± 40
11496	2113 ± 2	2115 ± 2	160 ± 5	167 ± 4	20 ± 14	44 ± 16	170 ± 85	96 ± 29
11498	3285 ± 5	3284 ± 4	251 ± 1	251 ± 2	71 ± 2	71 ± 2	268 ± 9	274 ± 9
11557	1393 ± 1	1392 ± 1	272 ± 2	276 ± 3	$29 \pm 22^*$	$29 \pm 22^*$	109 ± 6	105 ± 72
11597	40 ± 1	40 ± 2	241 ± 1	241 ± 3	41 ± 1	40 ± 10	155 ± 3	154 ± 32
11670	769 ± 1	776 ± 3	327 ± 2	333 ± 2	69 ± 2	65 ± 2	191 ± 8	190 ± 9
11707	900 ± 1	897 ± 2	58 ± 2	59 ± 3	63 ± 6	70 ± 4	91 ± 7	97 ± 5
11852	5822 ± 2	5821 ± 3	188 ± 2	189 ± 3	56 ± 6	47 ± 7	216 ± 14	221 ± 27
11861	1477 ± 1	1476 ± 2	214 ± 1	218 ± 3	49 ± 4	43 ± 12	176 ± 8	181 ± 39
11872	1138 ± 1	1140 ± 1	81 ± 1	86 ± 2	42 ± 4	47 ± 3	187 ± 12	183 ± 12
11891	468 ± 5	466 ± 6	114 ± 11	119 ± 10	$43 \pm 23^*$	$43 \pm 23^*$	85 ± 19	83 ± 35
11914	944 ± 1	945 ± 1	265 ± 1	266 ± 2	37 ± 2	33 ± 4	276 ± 10	285 ± 26
11951	1083 ± 1	1085 ± 2	258 ± 2	261 ± 4	75 ± 2	76 ± 8	100 ± 8	106 ± 7
12060	878 ± 1	879 ± 2	187 ± 2	187 ± 3	39 ± 4	36 ± 11	97 ± 5	107 ± 27
12082	793 ± 1	792 ± 2	137 ± 4	143 ± 5	30 ± 20	14 ± 19	51 ± 74	105 ± 137
12101	766 ± 1	770 ± 2	319 ± 3	317 ± 4	$58 \pm 9^*$	$58 \pm 9^*$	101 ± 8	94 ± 12
12212	896 ± 2	899 ± 2	101 ± 11	82 ± 7	$71 \pm 27^*$	$71 \pm 27^*$	168 ± 67	78 ± 15
12276	5644 ± 1	5642 ± 2	323 ± 4	322 ± 5	39 ± 16	33 ± 15	77 ± 66	94 ± 37
12343	2372 ± 1	2371 ± 2	204 ± 1	203 ± 2	48 ± 2	52 ± 4	225 ± 4	221 ± 14

(1) Name of the galaxy in the Uppsala General Catalogue, except for IC 476, IC 2542, and NGC 5296 which do not have a UGC name. (2) Systemic velocity derived by our fitting method. (3) Systemic velocity as found in Epinat et al. (2008b). (4) Position angle of the receding side derived by our fitting method. (5) Position angle of the receding side as found in Epinat et al. (2008b). Those marked with an asterisk (*) have been fixed to the morphological value from HyperLeda. (6) Disc inclination derived by our fitting method. (7) Disc inclination as found in Epinat et al. (2008b). Those marked with an asterisk (*) have been fixed to the morphological value from HyperLeda, except for UGC 6118, UGC 9013, UGC 9363, and UGC 10791 for which we used the inclinations derived from HI data (Epinat et al. 2008b). (8) Peak circular velocity derived by our fitting method. (9) Peak circular velocity as found in Epinat et al. (2008b).

Table C2. Kinematic properties of the DYNAMO sample (see Section 4.3).

Galaxy (DYNAMO ID) (1)	Seeing (arcsec) (2)	$r_{\rm exp,r}$ (kpc) (3)	$r_{\rm t}$ (kpc) (4)	$r_{\rm t,Green}$ (kpc) (5)	i (°) (6)	$V_{2.2{\rm R_r}}$ (km s^{-1}) (7)	$V_{2.2{\rm R_r,Green}}$ (km s^{-1}) (8)	$\sigma_{\rm model}$ (km s^{-1}) (9)	$\sigma_{\rm m,Green}$ (km s^{-1}) (10)
A 04-3	2.1	4.9	1.37 ± 0.08	0.5 ± 0.1	58	232 ± 3	218 ± 2	1	10
A 04-4	1.4	6.2	2.98 ± 0.17	2.0 ± 0.1	54	135 ± 4	119 ± 4	20	29
A 08-3	1.3	5.9	2.00 ± 0.10	1.8 ± 0.0	59	146 ± 2	149 ± 3	20	19
A 08-4	4.0	7.5	10.29 ± 0.40	11.9 ± 0.1	80	381 ± 11	336 ± 23	1	44
A 10-1	1.0	7.2	2.00 ± 0.19	1.8 ± 0.1	76	217 ± 6	217 ± 3	39	39
A 10-2	4.0	7.4	6.04 ± 0.35	12.4 ± 0.1	81	232 ± 12	264 ± 19	1	41
A 13-2	1.1	5.6	2.30 ± 0.22	2.6 ± 0.1	35	127 ± 6	128 ± 12	18	23
B 10-1	1.5	4.7	1.68 ± 0.12	1.9 ± 0.1	66	135 ± 3	138 ± 4	25	27
B 14-1	1.5	5.9	0.78 ± 0.07	0.4 ± 0.1	78	199 ± 2	197 ± 2	1	17
B 15-1	1.0	5.6	1.70 ± 0.10	1.8 ± 0.1	44	187 ± 3	185 ± 6	29	30
B 15-2	1.0	5.4	0.94 ± 0.12	1.0 ± 0.1	55	123 ± 2	124 ± 3	26	24
B 20-1	1.1	5.2	1.46 ± 0.12	1.6 ± 0.1	35	105 ± 2	105 ± 5	27	29
C 00-1	1.4	3.0	1.63 ± 0.12	2.0 ± 0.1	56	104 ± 3	107 ± 8	30	32
C 13-1	1.3	5.3	1.42 ± 0.10	1.6 ± 0.1	31	223 ± 3	223 ± 8	26	29
C 13-3	1.4	5.1	0.38 ± 0.10	0.5 ± 0.1	72	203 ± 2	220 ± 2	1	11
C 14-2	1.4	2.7	0.40 ± 0.08	0.6 ± 0.0	58	159 ± 2	164 ± 4	26	26
C 20-2	1.9	2.6	0.70 ± 0.08	0.6 ± 0.1	45	178 ± 3	170 ± 6	1	4
C 21-1	1.4	1.9	0.47 ± 0.06	0.5 ± 0.1	58	239 ± 2	246 ± 9	1	41
C 22-2	1.3	2.8	0.87 ± 0.09	1.0 ± 0.1	35	198 ± 3	197 ± 9	29	40
D 00-2	1.3	2.8	0.48 ± 0.10	0.8 ± 0.1	35	61 ± 2	62 ± 2	35	38
D 10-4	1.4	2.8	13.37 ± 0.05	13.4 ± 0.1	68	65 ± 2	59 ± 23	60	63
D 13-1	1.4	3.6	2.56 ± 0.13	3.4 ± 0.1	76	113 ± 3	120 ± 10	34	36
D 13-5	1.4	4.4	0.40 ± 0.05	0.5 ± 0.1	49	192 ± 1	192 ± 2	35	46
D 15-1	1.4	2.0	0.34 ± 0.10	0.4 ± 0.1	55	149 ± 5	142 ± 5	23	24
D 15-2	1.6	0.6	0.30 ± 0.05	0.3 ± 0.1	38	59 ± 25	42 ± 22	39	42
D 15-3	1.2	4.7	0.61 ± 0.06	0.6 ± 0.0	47	244 ± 1	240 ± 3	25	31
D 20-1	1.1	2.9	0.74 ± 0.07	0.9 ± 0.0	39	134 ± 2	137 ± 5	35	37
D 21-3	1.3	4.2	0.31 ± 0.05	0.3 ± 0.1	57	171 ± 1	175 ± 2	28	23
D 22-1	1.3	2.0	0.36 ± 0.05	0.4 ± 0.1	20	144 ± 1	141 ± 10	29	31
D 22-2	1.3	5.0	3.55 ± 0.11	3.8 ± 0.1	54	155 ± 2	155 ± 6	33	35
D 23-1	1.4	3.0	1.51 ± 0.09	2.0 ± 0.1	54	184 ± 3	186 ± 10	35	35
E 04-1	0.9	3.5	0.73 ± 0.07	0.9 ± 0.2	35	372 ± 3	395 ± 22	17	58
E 09-1	1.3	4.2	1.41 ± 0.18	1.1 ± 0.1	19	235 ± 8	267 ± 46	35	37
E 10-1	1.6	3.9	0.68 ± 0.08	0.8 ± 0.2	31	198 ± 2	231 ± 10	33	45
G 03-2	1.7	2.8	0.64 ± 0.06	0.8 ± 0.2	24	189 ± 2	174 ± 24	32	43
G 04-1	1.8	3.1	0.86 ± 0.12	0.8 ± 0.1	24	264 ± 6	269 ± 22	34	50
G 08-4	1.1	2.6	0.70 ± 0.10	1.1 ± 0.1	30	104 ± 3	114 ± 11	41	49
G 10-1	2.1	2.5	0.72 ± 0.11	4.3 ± 0.5	49	117 ± 3	103 ± 17	52	62
G 11-1	1.8	4.3	0.67 ± 0.09	2.0 ± 0.2	23	302 ± 5	314 ± 37	39	47
G 13-1	1.5	3.2	0.68 ± 0.06	0.8 ± 0.2	69	103 ± 1	112 ± 2	68	76
G 14-1	1.1	2.5	0.65 ± 0.06	0.8 ± 0.2	35	147 ± 3	136 ± 8	64	70
G 14-3	1.8	1.6	0.74 ± 0.12	2.7 ± 0.3	41	254 ± 11	159 ± 40	52	81
G 20-2	0.9	2.0	0.69 ± 0.05	0.8 ± 0.2	34	173 ± 1	166 ± 10	37	45

(1) Name of the galaxy in the DYNAMO catalogue. (2) The full width at half-maximum seeing used in the fitting procedure. (3) Exponential disc scale radius from SDSS r-band photometry. (4) Turn-over radius derived by our fitting method. (5) Turn-over radius of Green et al. (2014) after applying a correction factor of 0.7 (for details, see Section 4.3). (6) Disc inclination from SDSS r-band photometry. (7) Circular velocity at a radius of 2.2 × r-band exponential disc scale lengths derived by our fitting method. (8) Circular velocity at a radius of 2.2 × r-band exponential disc scale lengths as found in Green et al. (2014). (9) Velocity dispersion derived by our fitting method. The statistical error is 1 – 3 km s^{-1}. (10) Velocity dispersion as found in Green et al. (2014).